arXiv:1507.08613v4 [stat.CO] 3 Feb 2017

Local Likelihood Estimation for Covariance Functions with Spatially-Varying Parameters: The convoSPAT Package for R

Mark D. Risser
The Ohio State University

Catherine A. Calder
The Ohio State University

Abstract

In spite of the interest in and appeal of convolution-based approaches for nonstationary spatial modeling, off-the-shelf software for model fitting does not as of yet exist. Convolution-based models are highly flexible yet notoriously difficult to fit, even with relatively small data sets. The general lack of pre-packaged options for model fitting makes it difficult to compare new methodology in nonstationary modeling with other existing methods, and as a result most new models are simply compared to stationary models. Using a convolution-based approach, we present a new nonstationary covariance function for spatial Gaussian process models that allows for efficient computing in two ways: first, by representing the spatially-varying parameters via a discrete mixture or "mixture component" model, and second, by estimating the mixture component parameters through a local likelihood approach. In order to make computations for a convolution-based nonstationary spatial model readily available, this paper also presents and describes the **convoSPAT** package for R. The nonstationary model is fit to both a synthetic data set and a real data application involving annual precipitation to demonstrate the capabilities of the package.

Keywords: spatial statistics, nonstationary modeling, local likelihood estimation, precipitation, R.

1. Introduction

The Gaussian process is an extremely popular modeling approach in modern-day spatial and environmental statistics, due largely to the fact that the model is completely characterized by first- and second-order properties, and the second-order properties are straightforward to specify through widely used classes of valid covariance functions. A broad literature on covariance function modeling exists, but traditional approaches are mostly based on assumptions of isotropy or stationarity, in which the covariance between the spatial process at two locations is a function of only the separation distance or separation vector, respectively. This modeling assumption is made mostly for convenience, and is rarely a realistic assumption in practice. As a result, a wide variety of nonstationary covariance function models for Gaussian process models have been developed (e.g., Sampson and Guttorp 1992; Higdon 1998; Damian, Sampson, and Guttorp 2001; Fuentes 2001; Schmidt and O'Hagan 2003; Paciorek and Schervish 2006; Calder 2008; Schmidt, Guttorp, and O'Hagan 2011; Reich, Eidsvik, Guindani, Nail, and Schmidt 2011; and Vianna Neto, Schmidt, and Guttorp 2014), in which

the spatial dependence structure is allowed to vary over the spatial region of interest. However, while these nonstationary approaches more appropriately model the covariance in the spatial process, most are also highly complex and require intricate model-fitting algorithms, making it very difficult to replicate their results in a general setting. Therefore, when new nonstationary methods are developed, their performance is usually compared to stationary models, for which robust software is available. While software exists for several of these nonstationary approaches (see below), there are currently no pre-packaged options for fitting convolution-based nonstationary models.

To address this need, we present a simplified version of the nonstationary spatial Gaussian process model introduced by Paciorek and Schervish (2006) in which the locally-varying geometric anisotropies are modeled using a "mixture component" approach, similar to the discrete mixture kernel convolution approach in Higdon (1998) but allowing the underlying correlation structure to be specified by the modeler. The model is extended to allow other properties to vary over space as well, such as the process variance, nugget effect, and smoothness. An additional degree of efficiency is gained by using local likelihood techniques to estimate the spatially-varying features of the spatial process; then, the locally estimated features are smoothed over space, similar in nature to the approach of Fuentes (2002).

This paper also presents and describes the **convoSPAT** package for R for conducting a full analysis of point-referenced spatial data using a convolution-based nonstationary spatial Gaussian process model. The primary contribution of the package is to provide accessible model-fitting tools for spatial data of this type, as software for convolution-based nonstationary modeling does not currently exist. Furthermore, the methods used by the package are computationally efficient even when the size of the data is relatively large (on the order of $n = 1000$). The package is able to handle both a single realization of the spatial process observed at a finite set of locations, as well as independent and identically distributed replicates of the spatial process observed at a common set of locations. Finally, the paper demonstrates how the package can be used, and provides analyses of both simulated and real data sets.

As noted, there are several other (albeit non convolution-based) methods for nonstationary spatial modeling that do offer software, namely the basis function approach in the **fields** package (Nychka, Furrer, and Sain 2014) and the Gaussian Markov random field approach in the **INLA** package (Lindgren, Rue, and Lindstrom 2011; Ingebrigtsen, Lindgren, and Steinsland 2014; Fuglstad, Lindgren, Simpson, and Rue 2015; Lindgren and Rue 2015), both available for R. However, as these methods arise from significantly different modeling approaches, the **convoSPAT** package represents a novel contribution to the set of available software for nonstationary spatial modeling. Comparison across software for these various packages is beyond the scope of this paper and will be reserved for future work.

The paper is organized as follows. Section 2 introduces a convolution-based approach for nonstationary spatial statistical modeling, and Section 3 describes a full model for observed data and the mixture component parameterization. Section 4 outlines a computationally efficient approach to inference for the model introduced in Section 3; Sections 5, 6, and 7 outline usage of the **convoSPAT** package and present two applications. Section 8 concludes the paper.

2. A convolution-based nonstationary covariance function

Process convolutions or moving average models are popular constructive methods for specifying a nonstationary process model. In general, a spatial stochastic process $Y(\cdot)$ on $G \subset \mathcal{R}^d$ can be defined by the kernel convolution

$$Y(\mathbf{s}) = \int_{\mathcal{R}^d} K_{\mathbf{s}}(\mathbf{u}) dW(\mathbf{u}), \tag{1}$$

where $W(\cdot)$ is a d-dimensional stochastic process and $K_{\mathbf{s}}(\cdot)$ is a (possibly parametric) spatially-varying kernel function centered at $\mathbf{s} \in G$. Higdon (2002) summarizes the extremely flexible class of spatial process models defined by Equation 1: see, for example, Barry and Ver Hoef (1996), Ver Hoef, Cressie, and Barry (2004), Ickstadt and Wolpert (1998), Higdon (1998), Ver Hoef *et al.* (2004), and Hoef and Barry (1998).

The kernel convolution Equation 1 defines a mean-zero nonstationary spatial Gaussian process (GP) if $W(\cdot)$ is chosen to be d-dimensional Brownian motion. A benefit of using Equation 1 is that in this case the kernel functions completely specify the second-order properties of the GP through the covariance function

$$\text{Cov}(Y(\mathbf{s}), Y(\mathbf{s}')) = E\big[Y(\mathbf{s})Y(\mathbf{s}')\big] = \int_G K_{\mathbf{s}}(\mathbf{u}) K_{\mathbf{s}'}(\mathbf{u}) d\mathbf{u}, \tag{2}$$

where $\mathbf{s}, \mathbf{s}' \in G$. The popularity of this approach is due largely to the fact that it is much easier to specify kernel functions than a covariance function directly, since the kernel functions only require $\int_{\mathcal{R}^d} K_{\mathbf{s}}(\mathbf{u}) d\mathbf{u} < \infty$ and $\int_{\mathcal{R}^d} K^2_{\mathbf{s}}(\mathbf{u}) d\mathbf{u} < \infty$, while a covariance function must be even and nonnegative definite (Bochner 1959; Adler 1981). A famous result (Thiébaux 1976; Thiébaux and Pedder 1987) uses a parametric class of Gaussian kernel functions in Equation 2 to give a closed-form covariance function; this result was later extended (Paciorek 2003; Paciorek and Schervish 2006; Stein 2005) to show that

$$C^{NS}(\mathbf{s}, \mathbf{s}'; \boldsymbol{\theta}) = \sigma(\mathbf{s})\sigma(\mathbf{s}') \frac{|\boldsymbol{\Sigma}(\mathbf{s})|^{1/4} |\boldsymbol{\Sigma}(\mathbf{s}')|^{1/4}}{\left| \frac{\boldsymbol{\Sigma}(\mathbf{s}) + \boldsymbol{\Sigma}(\mathbf{s}')}{2} \right|^{1/2}} g\left(\sqrt{Q(\mathbf{s}, \mathbf{s}')} \right), \tag{3}$$

is a valid, nonstationary, parametric covariance function on $\mathcal{R}^d, d \geq 1$, when $g(\cdot)$ is chosen to be a valid correlation function on $\mathcal{R}^d, d \geq 1$. Note that Equation 3 no longer requires kernel functions to be specified. In Equation 3, $\boldsymbol{\theta}$ is a generic parameter vector, $\sigma(\cdot)$ represents a spatially-varying standard deviation, $\boldsymbol{\Sigma}(\cdot)$ is a $d \times d$ matrix that represents the spatially-varying local anisotropy (controlling both the range and direction of dependence), and

$$Q(\mathbf{s}, \mathbf{s}') = (\mathbf{s} - \mathbf{s}')^\top \left(\frac{\boldsymbol{\Sigma}(\mathbf{s}) + \boldsymbol{\Sigma}(\mathbf{s}')}{2} \right)^{-1} (\mathbf{s} - \mathbf{s}') \tag{4}$$

is a Mahalanobis distance. Furthermore, choosing $g(\cdot)$ to be the Matérn correlation function also allows for the introduction of $\kappa(\cdot)$, a spatially-varying smoothness parameter (Stein 2005; in this case, the Matérn correlation function in Equation 3 has smoothness $[\kappa(\mathbf{s}) + \kappa(\mathbf{s}')]/2$). While using Equation 3 no longer requires the notion of kernel convolution, we refer to $\boldsymbol{\Sigma}(\cdot)$ as the kernel matrix, since it was originally defined as the covariance matrix of a Gaussian kernel function (Thiébaux 1976; Thiébaux and Pedder 1987). The covariance function in Equation 3 is extremely flexible, and has been used in various forms throughout the literature, e.g., Paciorek and Schervish (2006), Anderes and Stein (2011), Kleiber and Nychka (2012), Katzfuss (2013), and Risser and Calder (2015).

3. A nonstationary spatial Gaussian process model

The covariance function in Equation 3 can be used to define a nonstationary spatial Gaussian process model using the following framework. Let $\{Z(\mathbf{s}), \mathbf{s} \in G\}$ be a spatial field defined on $G \subset \mathcal{R}^d$, where

$$Z(\mathbf{s}) = \mathbf{x}(\mathbf{s})^\top \boldsymbol{\beta} + Y(\mathbf{s}) + \epsilon(\mathbf{s}). \tag{5}$$

In Equation 5, the mean of the spatial field is $E[Z(\mathbf{s})] = \mathbf{x}(\mathbf{s})^\top \boldsymbol{\beta}$, where $\mathbf{x}(\mathbf{s})$ is a p-vector of covariates for location $\mathbf{s}$ and $\boldsymbol{\beta} \in \mathcal{R}^p$ are unknown regression coefficients. $Y(\cdot)$ represents a spatially-dependent, mean-zero Gaussian process with covariance function C^{NS} from Equation 3, while $\epsilon(\cdot)$ represents measurement error and, given $\tau^2(\cdot)$, is conditionally independent $\mathcal{N}(0, \tau^2(\mathbf{s}))$. (Note: $\mathcal{N}(a, b)$ denotes the univariate Gaussian distribution with mean a and variance b.) The spatially-referenced random components, $\epsilon(\cdot)$ and $Y(\cdot)$, are assumed to be mutually independent. Finally, define $\boldsymbol{\theta}$ to be a vector of all the variance-covariance parameters from the Gaussian process $Y(\cdot)$ and error process $\epsilon(\cdot)$.

Now, suppose we have observations which are a partial realization of $Z(\cdot)$, taken at a fixed, finite set of n spatial locations $\{\mathbf{s}_i\}_{i=1}^n \in G$, giving the random (observed) vector $\mathbf{Z} = (Z(\mathbf{s}_1), ..., Z(\mathbf{s}_n))$. The model in Equation 5 implies that $\mathbf{Z}$ has a multivariate Gaussian distribution, conditional on the unobserved latent process $\mathbf{Y} = (Y(\mathbf{s}_1), ..., Y(\mathbf{s}_n))$ and all other model parameters:

$$\mathbf{Z}|\mathbf{Y}, \boldsymbol{\beta}, \boldsymbol{\theta} \sim \mathcal{N}_n\big(\mathbf{X}\boldsymbol{\beta} + \mathbf{Y}, \mathbf{D}(\boldsymbol{\theta})\big), \tag{6}$$

where the ith row of $\mathbf{X}$ is $\mathbf{x}(\mathbf{s}_i)$ and $\mathbf{D}(\boldsymbol{\theta})$ is a diagonal matrix with (i, i) element $\tau^2(\mathbf{s}_i)$. (Note: $\mathcal{N}_q(\mathbf{a}, \mathbf{B})$ denotes the q-variate Gaussian distribution with mean vector $\mathbf{a}$ and covariance $\mathbf{B}$.) Integrating out the process $\mathbf{Y}$ from Equation 6, we can obtain the marginal likelihood of the observed data $\mathbf{Z}$ given all parameters, which is

$$\mathbf{Z}|\boldsymbol{\beta}, \boldsymbol{\theta} \sim \mathcal{N}_n\big(\mathbf{X}\boldsymbol{\beta}, \mathbf{D}(\boldsymbol{\theta}) + \boldsymbol{\Omega}(\boldsymbol{\theta})\big), \tag{7}$$

where $\boldsymbol{\Omega}(\boldsymbol{\theta})$ has elements $\Omega_{ij}(\boldsymbol{\theta}) = C^{NS}(\mathbf{s}_i, \mathbf{s}_j; \boldsymbol{\theta})$. For a particular application, the practitioner can specify the underlying correlation structure (through $g(\cdot)$) as well as determine which of $\{\boldsymbol{\Sigma}(\cdot), \sigma(\cdot), \tau^2(\cdot)\}$ (or $\kappa(\cdot)$, if the Matérn is used) should be fixed or allowed to vary spatially. However, some care should be taken in choosing which quantities should be spatially-varying: for example, Anderes and Stein (2011) note that allowing both $\boldsymbol{\Sigma}(\cdot)$ and $\kappa(\cdot)$ to vary over space leads to issues with identifiability.

3.1. Discrete mixture representation

One way to reduce the computational demands of fitting a Gaussian process-based spatial model with parametric covariance function given by Equation 3 is by characterizing the nonstationary behavior of a spatial process through the discretized basis kernel approach of Higdon (1998). Higdon (1998) estimated the Gaussian kernel function for a generic location to be a weighted average of "basis" kernel functions, estimated locally over the spatial region of interest. However, since the use of Gaussian kernel functions results in undesirable smoothness properties (see, e.g., Paciorek and Schervish 2006), we instead use a related "mixture component" approach, in which the parametric quantities for an arbitrary spatial location are defined as a mixture of spatially-varying parameter values associated with a fixed set of component locations. Specifically, in this new approach, we define mixture component locations $\{\mathbf{b}_k : k = 1, \ldots, K\}$ with corresponding parameters $\{(\boldsymbol{\Sigma}_k, \sigma_k^2, \tau_k^2, \kappa_k) : k = 1, \ldots, K\}$ (which

are the kernel matrix, variance, nugget variance, and smoothness, respectively). Then, for $\phi \in \{\boldsymbol{\Sigma}, \sigma^2, \tau^2, \kappa\}$, the parameter set for an arbitrary location $\mathbf{s} \in G$ is calculated as

$$\phi(\mathbf{s}) = \sum_{k=1}^{K} w_k(\mathbf{s})\phi_k, \tag{8}$$

where

$$w_k(\mathbf{s}) \propto \exp\left\{-\frac{||\mathbf{s} - \mathbf{b}_k||^2}{2\lambda_w}\right\} \tag{9}$$

such that $\sum_{k=1}^{K} w_k(\mathbf{s}) = 1$. For example, the kernel matrix for $\mathbf{s} \in G$ is $\boldsymbol{\Sigma}(\mathbf{s}) = \sum_{k=1}^{K} w_k(\mathbf{s})\boldsymbol{\Sigma}_k$. In Equation 9, λ_w acts as a tuning parameter, ensuring that the rate of decay in the weighting function is appropriate for both the data set and scale of the spatial domain. Using this approach, the number of parameters is now linear in K, the number of mixture component locations, instead of n, the sample size. Furthermore, this specification still enables the modeler to choose which parameters should be spatially-varying: the kernel matrices, the process variance, the nugget variance, and the smoothness.

3.2. Prediction

Define $\mathbf{Z}^* = (Z(\mathbf{s}_1^*), ..., Z(\mathbf{s}_m^*))$ to be a vector of the process values at all prediction locations of interest. The Gaussian process model in Equation 5 implies that

$$\left[\begin{array}{c|c} \mathbf{Z} \\ \mathbf{Z}^* \end{array} \middle| \boldsymbol{\beta}, \boldsymbol{\theta} \right] \sim \mathcal{N}_{n+m}\left(\left[\begin{array}{c} \mathbf{X}\boldsymbol{\beta} \\ \mathbf{X}^*\boldsymbol{\beta} \end{array}\right], \left[\begin{array}{cc} \mathbf{D}(\boldsymbol{\theta}) + \boldsymbol{\Omega}(\boldsymbol{\theta}) & \boldsymbol{\Omega}_{\mathbf{Z}\mathbf{Z}^*}(\boldsymbol{\theta}) \\ \boldsymbol{\Omega}_{\mathbf{Z}^*\mathbf{Z}}(\boldsymbol{\theta}) & \mathbf{D}^*(\boldsymbol{\theta}) + \boldsymbol{\Omega}^*(\boldsymbol{\theta}) \end{array}\right]\right),$$

where $\text{Cov}(\mathbf{Z}^*) = \mathbf{D}^*(\boldsymbol{\theta}) + \boldsymbol{\Omega}^*(\boldsymbol{\theta})$ and $\text{Cov}(\mathbf{Z}, \mathbf{Z}^*) = \boldsymbol{\Omega}_{\mathbf{Z}\mathbf{Z}^*}(\boldsymbol{\theta})$. By the properties of the multivariate Gaussian distribution,

$$\mathbf{Z}^*|\mathbf{Z} = \mathbf{z}, \boldsymbol{\beta}, \boldsymbol{\theta} \sim \mathcal{N}_m(\boldsymbol{\mu}_{\mathbf{Z}^*|\mathbf{z}}, \boldsymbol{\Sigma}_{\mathbf{Z}^*|\mathbf{z}}), \tag{10}$$

where

$$\boldsymbol{\mu}_{\mathbf{Z}^*|\mathbf{z}} = \mathbf{X}^*\boldsymbol{\beta} + \boldsymbol{\Omega}_{\mathbf{Z}^*\mathbf{Z}}(\boldsymbol{\theta})[\mathbf{D}(\boldsymbol{\theta}) + \boldsymbol{\Omega}(\boldsymbol{\theta})]^{-1}(\mathbf{z} - \mathbf{X}\boldsymbol{\beta}), \tag{11}$$

and

$$\boldsymbol{\Sigma}_{\mathbf{Z}^*|\mathbf{z}} = [\mathbf{D}^*(\boldsymbol{\theta}) + \boldsymbol{\Omega}^*(\boldsymbol{\theta})] - \boldsymbol{\Omega}_{\mathbf{Z}^*\mathbf{Z}}(\boldsymbol{\theta})[\mathbf{D}(\boldsymbol{\theta}) + \boldsymbol{\Omega}(\boldsymbol{\theta})]^{-1}\boldsymbol{\Omega}_{\mathbf{Z}\mathbf{Z}^*}(\boldsymbol{\theta}). \tag{12}$$

Using plug-in estimates $\widehat{\boldsymbol{\beta}}$ and $\widehat{\boldsymbol{\theta}}$ (see Section 4), the predictor for $\mathbf{Z}^*$ is then $\widehat{\boldsymbol{\mu}}_{\mathbf{Z}^*|\mathbf{z}}$ with corresponding prediction errors as the square root of the diagonal elements of $\widehat{\boldsymbol{\Sigma}}_{\mathbf{Z}^*|\mathbf{z}}$.

Out-of-sample evaluation criteria

Three cross-validation evaluation criteria can be used to assess the fit of the nonstationary spatial model given in Equation 5. First, the mean squared prediction error

$$MSPE = \frac{1}{m}\sum_{j=1}^{m}(z_j^* - \widehat{z}_j^*)^2, \tag{13}$$

where z_j^* is the jth held-out observed (or "validation") value and $\widehat{z}_j^*$ is the corresponding predictor (from Equation 11). The MSPE is a point-wise measure of model fit, and smaller MSPE indicates better predictions.

Second, to assess the prediction error relative to the standard error of each prediction, we use the so-called prediction mean squared deviation ratio

$$pMSDR = \frac{1}{m}\sum_{j=1}^{m}\frac{(z_j^* - \widehat{z}_j^*)^2}{\widehat{\sigma}_j}, \tag{14}$$

where z_j^* and $\widehat{z}_j^*$ are defined as above and $\widehat{\sigma}_j$ is the prediction error corresponding to $\widehat{z}_j^*$ (from Equation 12).

Finally, the continuous rank probability score will be used (a proper scoring rule; see Gneiting and Raftery 2007). For the jth prediction, this is defined as

$$CRPS_j \equiv CRPS(F_j, z_j^*) = -\int_{-\infty}^{\infty}\left(F_j(x) - 1\{x \geq z_j^*\}\right)^2 dx, \tag{15}$$

where $F_j(\cdot)$ is the cumulative distribution function (CDF) for the predictive distribution of z_j^* given the training data and $1\{\cdot\}$ is the indicator function. In this case, given that the predictive CDF is Gaussian (conditional on parameters; see Equation 10), a computational shortcut can be used for calculating Equation 15: when F is Gaussian with mean μ and variance σ^2,

$$CRPS(F, z_j^*) = \sigma\left[\frac{1}{\sqrt{\pi}} - 2\cdot\phi\left(\frac{z_j^* - \mu}{\sigma}\right) - \frac{z_j^* - \mu}{\sigma}\left(2\cdot\Phi\left(\frac{z_j^* - \mu}{\sigma}\right) - 1\right)\right],$$

where ϕ and Φ denote the probability density and cumulative distribution functions, respectively, of a standard Gaussian random variable. The reported metric will be the average over all validation locations, $\widehat{CRPS} = m^{-1}\sum_{j=1}^{m}\widehat{CRPS}_j$. CRPS measures the fit of the predictive density; larger CRPS (i.e., smaller negative values) indicates better model fit.

4. Computationally efficient inference

As discussed in Section 1, fast and efficient inference for a nonstationary process convolution model has yet to be made readily available for general use. In spite of its popularity, the use of Equation 3 always requires some kind of constraints and has suffered from a lack of widespread use due to the complexity of the requisite model fitting and limited pre-packaged options. Focusing on the spatially-varying local anisotropy matrices $\boldsymbol{\Sigma}(\cdot)$, the covariance function in Equation 3 requires a kernel matrix at every observation and prediction location of interest. Paciorek and Schervish (2006) accomplish this by modeling $\boldsymbol{\Sigma}(\cdot)$ as itself a (stationary) stochastic process, assigning Gaussian process priors to the elements of the spectral decomposition of $\boldsymbol{\Sigma}(\cdot)$; alternatively, Katzfuss (2013) uses a basis function representation of $\boldsymbol{\Sigma}(\cdot)$. Both of these models are highly parameterized and require intricate Markov chain Monte Carlo methods for model fitting.

The approach we propose provides efficiency in two ways: first, from the model itself, which uses a discrete mixture representation (see Section 3.1), and second, by fitting the mixture components of the model locally, using the idea of local likelihood estimation (Tibshirani and Hastie 1987).

4.1. Local likelihood estimation

Using the discrete mixture representation of Equation 8, a "full likelihood" approach to parameter estimation could be taken, in either a Bayesian or maximum likelihood framework, although the optimization in a maximum likelihood approach could become intractable for either moderately large K or large n. However, since the primary goal of this new methodology is computational speed, a further degree of efficiency can be gained by using local likelihood estimation (LLE; Tibshirani and Hastie 1987).

Before discussing the local likelihood approach, we outline a restricted maximum likelihood (REML; see Patterson and Thompson 1971, Patterson and Thompson 1974, and Kitanidis 1983) approach for separating estimation of the mean parameters $\boldsymbol{\beta}$ from the covariance parameters $\boldsymbol{\theta}$. The full log-likelihood for $\boldsymbol{\beta}$ and $\boldsymbol{\theta}$ in Equation 5 is

$$\mathcal{L}^F(\boldsymbol{\beta}, \boldsymbol{\theta}; \mathbf{Z}) = -\frac{1}{2}\log|\boldsymbol{\Omega} + \mathbf{D}| - \frac{1}{2}(\mathbf{Z} - \mathbf{X}\boldsymbol{\beta})^\top(\boldsymbol{\Omega} + \mathbf{D})^{-1}(\mathbf{Z} - \mathbf{X}\boldsymbol{\beta}), \tag{16}$$

where we have abbreviated $\mathbf{D} \equiv \mathbf{D}(\boldsymbol{\theta})$ and $\boldsymbol{\Omega} \equiv \boldsymbol{\Omega}(\boldsymbol{\theta})$; a standard maximum likelihood approach would set out to maximize $\mathcal{L}^F(\boldsymbol{\beta}, \boldsymbol{\theta}; \mathbf{Z})$ directly. REML, on the other hand, uses a (log) likelihood based on $n - p$ linearly independent linear combinations of the data that have an expected value of zero for all possible $\boldsymbol{\beta}$ and $\boldsymbol{\theta}$. Regardless of which set of linearly independent combinations is chosen, the "restricted" log-likelihood, which depends only on $\boldsymbol{\theta}$, is

$$\mathcal{L}^R(\boldsymbol{\theta}; \mathbf{Z}) = -\frac{1}{2}\log|\boldsymbol{\Omega} + \mathbf{D}| - \frac{1}{2}\log|\mathbf{X}^\top(\boldsymbol{\Omega} + \mathbf{D})^{-1}\mathbf{X}| - \frac{1}{2}\mathbf{Z}^\top\mathbf{P}\mathbf{Z}, \tag{17}$$

where

$$\mathbf{P} = (\boldsymbol{\Omega} + \mathbf{D})^{-1} - (\boldsymbol{\Omega} + \mathbf{D})^{-1}\mathbf{X}\big(\mathbf{X}^\top(\boldsymbol{\Omega} + \mathbf{D})^{-1}\mathbf{X}\big)^{-1}\mathbf{X}^\top(\boldsymbol{\Omega} + \mathbf{D})^{-1}. \tag{18}$$

The REML estimate of $\boldsymbol{\theta}$ is obtained by maximizing $\mathcal{L}^R(\boldsymbol{\theta}; \mathbf{Z})$, and the estimate of $\boldsymbol{\beta}$ is the generalized least squares estimate

$$\widehat{\boldsymbol{\beta}} = \big(\mathbf{X}^\top(\widehat{\boldsymbol{\Omega}} + \widehat{\mathbf{D}})^{-1}\mathbf{X}\big)^{-1}\mathbf{X}^\top(\widehat{\boldsymbol{\Omega}} + \widehat{\mathbf{D}})^{-1}\mathbf{Z}, \tag{19}$$

which is obtained by plugging in $\widehat{\boldsymbol{\theta}}$ to calculate $\widehat{\boldsymbol{\Omega}}$ and $\widehat{\mathbf{D}}$. These parameter estimates can then be plugged in to $\widehat{\boldsymbol{\mu}}_{\mathbf{Z}^*|\mathbf{Z}}$ and $\widehat{\boldsymbol{\Sigma}}_{\mathbf{Z}^*|\mathbf{Z}}$ to obtain predictions and prediction standard errors.

In the LLE approach, instead of maximizing Equation 17 directly we will set out to maximize $\mathcal{L}_k^R(\boldsymbol{\theta}_{N_k}; \mathbf{Z}_{N_k})$, where $N_k \equiv N_k(r)$ is a neighborhood for each mixture component location $\mathbf{b}_k$ that depends on the radius r, such that

$$N_k = \big\{\mathbf{s}_i \in \{\mathbf{s}_1, \ldots, \mathbf{s}_n\} : \{||\mathbf{s}_i - \mathbf{b}_k|| \leq r\}\big\}$$

and

$$\mathbf{Z}_{N_k} = \big\{Z(\mathbf{s}) : \mathbf{s} \in N_k\big\}.$$

Correspondingly, $\boldsymbol{\theta}_{N_k} = (\boldsymbol{\Sigma}_k, \sigma_k^2, \tau_k^2, \kappa_k)$. The radius r defines the "span" (Tibshirani and Hastie 1987) or window size for each mixture component. The restricted log-likelihood for neighborhood N_k will be based on a stationary version of the spatial model in Equation 5, namely

$$\widetilde{Z}(\mathbf{s}) = \mathbf{x}(\mathbf{s})^\top\widetilde{\boldsymbol{\beta}} + \widetilde{Y}(\mathbf{s}) + \widetilde{\epsilon}(\mathbf{s}), \tag{20}$$

where $\widetilde{Y}(\cdot)$ is a stationary, mean-zero spatial process with covariance function

$$C^S(\mathbf{s} - \mathbf{s}') = \sigma^2 g\left(||\boldsymbol{\Sigma}^{-1/2}(\mathbf{s} - \mathbf{s}')||\right), \tag{21}$$

the $\widetilde{\epsilon}(\cdot)$ are independent and identically distributed as $\mathcal{N}(0, \tau^2)$, conditional on τ^2, and again $\widetilde{Y}(\cdot)$ and $\widetilde{\epsilon}(\cdot)$ are independent. Again, in a REML framework, only the variance and covariance parameters $\{\boldsymbol{\Sigma}_k, \sigma_k^2, \tau_k^2, \kappa_k\}$ need to be estimated for each $k = 1, \ldots, K$. No estimates will be obtained for the local mean coefficient vector $\widetilde{\boldsymbol{\beta}}$, as all of the mean parameters will be estimated globally.

One final note regarding the estimation of the kernel matrices: the kernel matrix for mixture component location k will be obtained through estimating the parameters of its spectral decomposition, namely λ_1, λ_2, and η, where

$$\boldsymbol{\Sigma} = \begin{bmatrix} \cos(\eta) & -\sin(\eta) \\ \sin(\eta) & \cos(\eta) \end{bmatrix} \begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix} \begin{bmatrix} \cos(\eta) & \sin(\eta) \\ -\sin(\eta) & \cos(\eta) \end{bmatrix} \tag{22}$$

(in the case that we have fixed $d = 2$). Here, λ_1 and λ_2 are eigenvalues and represent squared ranges (such that $\lambda_1 > 0$ and $\lambda_2 > 0$) and η represents an angle of rotation, constrained to lie between 0 and $\pi/2$ for identifiability purposes (Katzfuss 2013).

The full model in Equation 5 can be fit after plugging REML estimates $\{\widehat{\boldsymbol{\Sigma}}_k, \widehat{\sigma}_k^2, \widehat{\tau}_k^2, \widehat{\kappa}_k : k = 1, \ldots, K\}$ into the covariance function in Equation 3 using the discrete basis representation in Equation 8 to calculate the likelihood for the observed data. Variance quantities that are not specified to be spatially-varying can then be estimated again using REML with the spatially-varying components considered fixed. For example, if for a particular model only $\boldsymbol{\Sigma}(\cdot)$ is allowed to vary spatially and the smoothness is fixed, it remains to estimate the overall nugget τ^2 and variance σ^2. The restricted Gaussian likelihood for these parameters is then

$$\mathcal{L}^R(\sigma^2, \tau^2; \mathbf{Z}, \mathbf{R}) = -\frac{1}{2}\log\left|\sigma^2\mathbf{R} + \tau^2\mathbf{I}_n\right| - \frac{1}{2}\log\left|\mathbf{X}^\top(\sigma^2\mathbf{R} + \tau^2\mathbf{I}_n)^{-1}\mathbf{X}\right| - \frac{1}{2}\mathbf{Z}^\top\mathbf{P}\mathbf{Z},$$

where $\mathbf{R}$ is the correlation matrix, i.e., the matrix calculated using Equation 3 without the $\sigma(\cdot)$ terms, and $\mathbf{P}$ is defined as in Equation 18. Once all of the covariance parameters have been estimated, the estimate of $\boldsymbol{\beta}$ can be calculated as in Equation 19.

Using this model requires both the number and placement of mixture component locations $\{\mathbf{b}_k : k = 1, \ldots, K\}$, selecting which of the spatial dependence parameters should be fixed or allowed to vary spatially, the tuning parameter for the weighting function λ_w, and the fitting radius r. Parameter estimates for this model are likely to be sensitive to the choice of K and the placement of mixture component locations. Furthermore, Tibshirani and Hastie (1987) discuss the importance of choosing r, which specifies the "span size," suggesting that the model should be fit using a range of r values, and use a global criterion such as the maximized overall likelihood, cross-validation, or Akaike's Information Criterion to choose the final model. This strategy could either be implemented on a trial-and-error basis or in an automated scheme. Of course, regardless of the number and locations of the mixture component centroids, the radius r should be chosen such that a large enough number of data points are used to estimate a local stationary model.

While different in both motivation and nature, the model outlined above is related to the local likelihood method described in Anderes and Stein (2011), which ties together locally stationary models to estimate a globally nonstationary model. The model in Anderes and Stein (2011) involves optimizing a sum of weighted increments of local log-likelihoods, where the weights are estimated smoothly using a smoothing kernel. Alternatively, our approach estimates spatially-varying parameters locally using only a subset of the data, then fixing

the global parameters according to Equation 8. Both of these approaches avoid the lack-of-smoothness issues innate to other similar segmentation approaches, such as Fuentes (2001) or the *ad hoc* nonstationary kriging approach in Paciorek and Schervish (2006), which Anderes and Stein (2011) call "hard thresholding" local likelihood estimates. Like Anderes and Stein (2011), our approach avoids the problem of non-smooth local parameter estimates implicit to hard thresholding methods by using the mixture component representation.

We conclude this overview of our methodology with a note regarding the computational demands of fitting this model. Recall that Gaussian likelihood calculations such as Equation 16 typically involve inverting and calculating the determinant of $n \times n$ matrices, requiring $\mathcal{O}(n^3)$ calculations (in this case, for each iteration of the optimization procedure). The local likelihood approach involves inverting a collection of K $n_k \times n_k$ matrices, so that the model requires more like $\mathcal{O}(K\overline{n}^3)$ calculations, where $\overline{n} = K^{-1}\sum_k n_k$ is the average local sample size. This represents a significant reduction in both the required memory and CPU time, so long as $\overline{n} << n$. However, note that since the local models are fit independently of each other, if parallelization is utilized (see Section 8) then the computational time could be further reduced to $\mathcal{O}(\overline{n}^3)$.

5. Using the convoSPAT package for R

The **convoSPAT** package (version 1.1.1) is available CRAN, and can be installed as usual:

```
R> install.packages( "convoSPAT" )
R> library( "convoSPAT" )
```

All of the data sets in Sections 6 and 7 are included in the package. The **convoSPAT** package uses functionality from the **ellipse** (Murdoch and Chow 2013), **fields** (Nychka *et al.* 2014), **geoR** (Ribeiro Jr. and Diggle 2015), **MASS** (Venables and Ripley 2002), **plotrix** (Lemon 2006), **sp** (Bivand, Pebesma, and Gomez-Rubio 2013; Pebesma and Bivand 2005), and **StatMatch** (D'Orazio 2016) packages for R.

Two notes should be made before discussing the functionality of the package. First, while the methods described in Sections 2, 3, and 4 are valid for spatial coordinates in a general d-dimensional Euclidean space, the following implementation only allows for two-dimensional coordinates, with $d = 2$. Second, on a more technical note, the package is implemented using the S3-classes of R, and therefore contains package-specific `predict` and `plot` functionality.

5.1. Nonstationary model fitting

The primary components of the **convoSPAT** package are the `NSconvo_fit` and `predict.NSconvo` functions which fit the nonstationary model discussed in Section 3 and provide predictions, respectively. The `NSconvo_fit` function takes the following arguments (with defaults as given):

```
NSconvo_fit( geodata = NULL, sp.SPDF = NULL,
  coords = geodata$coords, data = geodata$data,
  cov.model = "exponential", mean.model = data ~ 1, mc.locations = NULL,
  N.mc = NULL, mc.kernels = NULL, fit.radius, lambda.w = NULL,
```

```
ns.nugget = FALSE, ns.variance = FALSE, local.pars.LB = NULL,
local.pars.UB = NULL, global.pars.LB = NULL, global.pars.UB = NULL,
local.ini.pars = NULL, global.ini.pars = NULL )
```

The spatial coordinates and response variable of interest may be input in several different ways: first, the function accepts a `geodata` object from the **geoR** package (Ribeiro Jr. and Diggle 2015), using the `geodata` input; second, a `SpatialPointsDataFrame` object from the **sp** package (Bivand *et al.* 2013; Pebesma and Bivand 2005), using the `sp.SPDF` input; finally, the coordinates and data may be entered directly using the `coords` and `data` inputs. As a result, the package is able to handle a variety of geographic coordinate reference systems; note, however, that distances between points are calculated using a Mahalanobis distance (see Equation 4).

Two other required inputs are the number of mixture component locations (`N.mc`) and the `fit.radius` (previously denoted r). The user may specify a covariance model from the **geoR** (Ribeiro Jr. and Diggle 2015) options `cauchy`, `matern`, `circular`, `cubic`, `gaussian`, `exponential`, `spherical`, or `wave`, as well as a mean model through the usual formula notation (a constant mean is the default). For most applications (and as an alternative to specifying `N.mc`), the user will want to specify the mixture component locations directly: the default is to create an evenly spaced grid over the spatial domain of interest, which may not be appropriate if the spatial domain is non-rectangular. The tuning parameter for the weighting function λ_w is defined by `lambda.w`. The default for λ_w is fixed to be the square of one-half of the minimum distance between mixture component locations, or $(0.5 \min\{||\mathbf{b}_k - \mathbf{b}_{k'}||\})^2$ (in order to ensure a default scaling appropriate for the resolution of the mixture component grid), but may also be specified by the user. The user may also specify if either the nugget variance or process variance is to be spatially-varying by setting either `ns.nugget = TRUE` or `ns.variance = TRUE` (or both). If the mixture component kernels themselves are pre-specified (e.g., based on expert opinion), these may also be passed into the function, which will greatly reduce computational time.

Note that if the data and coordinates are not specified as a `geodata` object, the `data` argument for this function can accommodate replicates. This might be of interest for applications similar to the ones in Sampson and Guttorp (1992), in which the replicates represent repeated observations over time that have been temporally detrended. In this case, the model will assume a constant spatial dependence structure over replicates (time) as well as the same mean function over replicates (that is, the locations must be constant across replicates; furthermore, the regression coefficients will be constant across replicates).

The optimization method used within `optim` for this package is `"L-BFGS-B"`, which allows for the specification of upper and lower bounds for each parameter with respect to the optimization search. The upper and lower bounds may be passed to the function via `local.pars.LB`, `local.pars.UB`, `global.pars.LB`, and `global.pars.UB`. The local limits require vectors of length five, with bounds for the local parameters $\lambda_1, \lambda_2, \tau^2$, σ^2, and κ, while the global limits require vectors of length three, with bounds for the global parameters τ^2, σ^2, and κ. Default values for these limits are as follows: for both the global and local parameter estimation, the lower bounds for $\lambda_1, \lambda_2, \sigma^2, \tau^2$, and κ are fixed at 1e-5; the upper bound for the smoothness κ will be fixed to 30. The upper bounds for the variance and kernel parameters, on the other hand, will be specific to the application: for the nugget variance (τ^2) and process variance (σ^2), the upper bound will be $4\widehat{\sigma}^2_{OLS}$ (where $\widehat{\sigma}^2_{OLS}$ is the error variance estimate from a stan-

dard ordinary least squares procedure); the upper bound for λ_1 and λ_2 will be one-fourth of the maximum interpoint distance between observation locations in the data set. The bounds for η are fixed at 0 and $\pi/2$.

Given that many calls to `optim` are made within `NSconvo_fit`, the function prints a message to notify the user if `optim` returns any errors. In test runs of the package, the most common warning (non-fatal) message encountered is `"ABNORMAL_TERMINATION_IN_LINSRCH"`, which seems to have no negative impact on the results of the optimization.

The final options in the `NSconvo_fit` function involve `local.ini.pars` and `global.ini.pars`, which specify the initial values used for the local and global calls of `optim`, respectively. As with the limits, `local.ini.pars` is a vector of length five, with initial values for the local parameters $\lambda_1, \lambda_2, \tau^2$, σ^2, and κ, while `global.ini.pars` is a vector of length three, with initial values for the global parameters τ^2, σ^2, and κ. The default for these inputs are as follows: $\lambda_{1,init} = \lambda_{2,init} = c/10$, where c is the maximum interpoint distance between observation locations, $\tau^2_{init} = 0.1\widehat{\sigma}^2_{OLS}$, $\sigma^2_{init} = 0.9\widehat{\sigma}^2_{OLS}$, and $\kappa_{init} = 1$.

When the `NSconvo_fit` function is called, the status of the model fitting will be printed on the screen. As the function fits the locally stationary models for each mixture component location, the function will print the mixture component and number of observations that are currently being used for estimation. After the local models have been fit for each mixture component location, a printed message will notify the user that the variance parameters are being estimated globally (if applicable).

A function which may be helpful before running `NSconvo_fit` is the `mc_N` function, which returns the number of observations which will be used to fit each local model for a particular set of mixture component locations and fit radius. This function may be helpful when selecting the fit radius, as the user may want to know how many data points will be used to fit the local model for a number of different fit radii. The inputs to the function are

```
mc_N( coords, mc.locations, fit.radius )
```

where `coords` are the observation locations for the full data set, `mc.locations` are the mixture component locations, and `fit.radius` is the fitting radius. This function is also automatically run inside the `NSconvo_fit` wrapper, printing a warning message if any of the local sample sizes are less than 5. In this case, the user should refine the mixture component grid or expand the fit radius.

After the model fitting has completed, `NSconvo_fit` returns a `NSconvo` object which can be passed to the `summary.NSconvo` function to quickly summarize the fitted model. Among other elements, a `NSconvo` object includes:

- `mc.kernels`, which contains the estimated kernel matrices for the mixture component locations,
- `mc.locations`, which contains the mixture component locations,
- `MLEs.save`, which includes a data frame of the locally-estimated stationary model parameters for each mixture component location,
- `kernel.ellipses`, which includes the estimated kernel ellipse for each location in `coords`,

`beta.GLS`, the generalized least squares estimate of $\boldsymbol{\beta}$,

`beta.cov`, the estimated covariance matrix of $\widehat{\boldsymbol{\beta}}$,

`tausq.est`, the estimate of the nugget variance – either a constant (if estimated globally) or a vector with the estimated nugget variance for each location in `coords`,

`sigmasq.est`, the estimate of the process variance – either a constant (if estimated globally) or a vector with the estimated process variance for each location in `coords`,

`kappa.MLE`, the global estimate of the smoothness (for `cauchy` or `matern`),

`Cov.mat` and `Cov.mat.inv`, the estimated covariance matrix for the data and its inverse (respectively), and

`Xmat`, the design matrix for the mean model.

The `NSconvo` object can also be passed to the `predict.NSconvo` function, which calculates predictions $\widehat{\boldsymbol{\mu}}_{\mathbf{Z}^*|\mathbf{z}}$ and prediction standard errors $\widehat{\boldsymbol{\Sigma}}_{\mathbf{Z}^*|\mathbf{z}}$. The `predict.NSconvo` function takes the following arguments:

```
predict.NSconvo( object, pred.coords, pred.covariates = NULL, ... )
```

The `object` is the output of `NSconvo_fit`, `pred.coords` is a matrix of the prediction locations of interest, `pred.covariates` is a matrix of covariates for the prediction locations (the intercept is added automatically), and `...` allows other options to be passed to the default R `predict` methods. Calculating the predictions when the dimension of `pred.coords` is large is computationally expensive, and a progress meter prints while the machine is working. The output from `predict.NSconvo` includes:

`pred.means`, which contains the kriging predictor for each prediction location, and

`pred.SDs`, which contains the corresponding prediction standard error.

5.2. Anisotropic model fitting

For the sake of comparison, the functions `Aniso_fit` and `predict.Aniso` are also provided, which fit the stationary (anisotropic) model with the covariance function given by Equation 21 to the full dataset. Note: the following functions have been coded from scratch and represent reimplementations of methodology that already exists in other packages, such as **geoR** (Ribeiro Jr. and Diggle 2015). In spite of being reimplementations, the functions are provided for convenience, as their syntax matches `NSconvo_fit` and can be quickly implemented in tandem with `NSconvo_fit`. However, also note that the version provided here has not been optimized nearly to the extent of the version in **geoR**, and therefore `Aniso_fit` will take significantly longer than a corresponding implementation of, say, `likfit` in **geoR**.

The `Aniso_fit` function takes the following arguments (with defaults as given):

```
Aniso_fit( geodata = NULL, sp.SPDF = NULL,
  coords = geodata$coords, data = geodata$data,
  cov.model = "exponential", mean.model = data ~ 1, local.pars.LB = NULL,
  local.pars.UB = NULL, local.ini.pars = NULL )
```

The inputs to this function are identical to the corresponding inputs to the nonstationary model fitting function, and the output is an `Aniso` object. Among other elements, an `Aniso` object includes:

`MLEs.save`, which includes a data frame of the locally-estimated stationary model parameters for each mixture component location,

`beta.GLS`, the generalized least squares estimate of $\boldsymbol{\beta}$,

`beta.cov`, the estimated covariance matrix of $\widehat{\boldsymbol{\beta}}$,

`aniso.pars`, the global estimate of the anisotropy parameters λ_1, λ_2, and η, which define the anisotropy matrix in Equation 22,

`aniso.mat`, which gives the global estimate of $\boldsymbol{\Sigma}$ from Equation 22 in matrix form,

`tausq.est`, the global estimate of the nugget variance,

`sigmasq.est`, the global estimate of the process variance,

`kappa.MLE`, the global estimate of the smoothness (for `cauchy` or `matern`), and

`Cov.mat` and `Cov.mat.inv`, the estimated covariance matrix for the data and its inverse (respectively)

Once the anisotropic model has been fit and the output stored, the fitted model object can be passed to the `predict.Aniso` function, which (similar to the nonstationary predict function) calculates predictions $\widehat{\boldsymbol{\mu}}_{\mathbf{Z}^*|\mathbf{z}}$ and prediction standard errors $\widehat{\boldsymbol{\Sigma}}_{\mathbf{Z}^*|\mathbf{z}}$. The arguments are again identical to the nonstationary predict function:

```
predict.Aniso( object, pred.coords, pred.covariates = NULL, ... )
```

The `object` is the output of `Aniso_fit`, `pred.coords` is a matrix of the prediction locations of interest, and `pred.covariates` is a matrix of covariates for the prediction locations (the intercept is added automatically). Similar to the nonstationary predict function, the output from `predict.Aniso` includes:

`pred.means`, which contains the kriging predictor for each prediction location, and

`pred.SDs`, which contains the corresponding prediction standard error.

5.3. Evaluation criteria and plotting functions

This package includes a function to quickly calculate the evaluation criteria described in Section 3.2, as well as functions to visualize various components of the nonstationary model.

First, the `evaluate_CV` function calculates the MSPE, from Equation 13, the pMSDR, from Equation 14, and the CRPS, from Equation 15. The function inputs are simply

```
evaluate_CV( holdout.data, pred.mean, pred.SDs )
```

where `holdout.data` is the held-out validation data and `pred.mean` and `pred.SDs` are output from one of the fit functions. Note that the user must perform the subsetting of the data. The output of `evaluate_CV` is simply the MSPE, pMSDR, and CRPS, averaged over all hold-out locations.

Next, plotting functions are provided to help visualize the output of either the stationary or nonstationary model. The first is `plot.NSconvo`:

```
plot.NSconvo( x, plot.ellipses = TRUE, fit.radius = NULL,
  aniso.mat = NULL, true.mc = NULL, ref.loc = NULL,
  all.pred.locs = NULL, grid = TRUE, true.col = 1,
  aniso.col = 4, ns.col = 2, plot.mc.locs = TRUE, ... )
```

which plots either the estimated anisotropy ellipses (`plot.ellipses = TRUE`) or the estimated correlation (`plot.ellipses = FALSE`). The `x` is a `NSconvo` object; additional options can be added to a plot of the estimated anisotropy ellipses by specifying the `fit.radius` that was used to fit the model, the ellipse for the stationary model `aniso.mat` estimated in `Aniso_fit`, and the true mixture component ellipses (if known). To plot the estimated correlation, a reference location (`ref.loc`) must be specified, as well as all of the prediction locations of interest (`all.pred.locs`) and whether or not the predictions lie on a rectangular grid. The other options correspond to color values for the true mixture component ellipses (`true.col`), the anisotropy ellipse (`aniso.col`), and the estimated mixture component ellipses (`ns.col`); `plot.mc.locs` indicates whether or not the mixture component locations should be plotted. Note that the plotted ellipse is the 0.5 probability ellipse for a bivariate Gaussian random variable with covariance equal to the kernel matrix.

Similarly, the `plot.Aniso` function is provided to plot the estimated correlations from the stationary model. The function is

```
plot.Aniso( x, ref.loc = NULL, all.pred.locs = NULL, grid = TRUE, ... )
```

The inputs are identical to the `plot.NSconvo` function, except that the `object` must be of the `Aniso` class.

5.4. Other functions

Two additional functions are also provided to simulate data from the nonstationary model discussed in Section 3, and are used to create the simulated data set in Section 6. First is `f_mc_kernels`, which calculates the true mixture component kernel matrices through a generalized linear model for each component of the kernel matrices' spectral decomposition as given in Equation 22. The function, with default settings, is

```
f_mc_kernels( y.min = 0, y.max = 5, x.min = 0, x.max = 5, N.mc = 3^2,
  lam1.coef = c( -1.3, 0.5, -0.6 ), lam2.coef = c( -1.4, -0.1, 0.2 ),
  logit.eta.coef = c( 0, -0.15, 0.15 ) )
```

The inputs `y.mon`, `y.max`, `x.min`, and `x.max` define a rectangular spatial domain, `N.mc` is the number of mixture component locations, and `lam1.coef`, `lam2.coef`, and `logit.eta.coef`

define regression coefficients for the spatially-varying parameters λ_1, λ_2, and η. For example, taking `lam1.coef` $\equiv \boldsymbol{\beta}_{\lambda_1} = (\beta_0^{\lambda_1}, \beta_1^{\lambda_1}, \beta_2^{\lambda_1})^\top$, the first eigenvector for an arbitrary location $\mathbf{s} = (s_1, s_2)$ is

$$\lambda_1(\mathbf{s}) = \exp\{\beta_0^{\lambda_1} + \beta_1^{\lambda_1} s_1 + \beta_2^{\lambda_1} s_2\}.$$

The other eigenvalue is calculated similarly. The angle of rotation, on the other hand, is calculated through the scaled inverse logit transformation

$$\eta(\mathbf{s}) = \frac{\pi}{2} \cdot \text{logit}^{-1}\big(\beta_0^{\eta} + \beta_1^{\eta} s_1 + \beta_2^{\eta} s_2\big),$$

where `logit.eta.coef` $\equiv \boldsymbol{\beta}_{\eta} = (\beta_0^{\eta}, \beta_1^{\eta}, \beta_2^{\eta})^\top$. The default coefficients are those used to generate the simulated data set in Section 6, and were obtained by trial and error. The output of this function includes

`mc.locations`, which contains the mixture component locations, and

`mc.kernels`, which contains the mixture component kernels for each mixture component location.

The true mixture component kernel matrices generated in `f_mc_kernels` can be used to simulate a data set using the function

```
NSconvo_sim( grid = TRUE, y.min = 0, y.max = 5, x.min = 0, x.max = 5,
  N.obs = 20^2, sim.locations = NULL, mc.kernels.obj = NULL,
  mc.kernels = NULL, mc.locations = NULL, tausq = 0.1, sigmasq = 1,
  beta.coefs = 4, kappa = NULL, covariates = rep( 1, N.obs ),
  cov.model = "exponential" )
```

In this function, `grid` is a logical input specifying if the simulated data should lie on a grid (`TRUE`) or not (`FALSE`), `y.min`, `y.max`, `x.min`, and `x.max` define the rectangular spatial domain, `N.obs` specifies the number of observed locations, `mc.kernels.obj` is an object from `f_mc_kernels`, `tausq`, `sigmasq`, `beta.coefs`, and `kappa` specify the true parameter values, `covariates` specifies the design matrix for the mean function, and `cov.model` specifies the covariance model. The output of this function is

`sim.locations`, which contains the simulated data locations,

`mc.locations`, which contains the mixture component locations,

`mc.kernels`, which contains the mixture component kernels for each mixture component location,

`kernel.ellipses`, which contains the kernel matrices for each simulated data location,

`Cov.mat`, which contains the true covariance matrix of the simulated data, and

`sim.data`, which contains the simulated data.

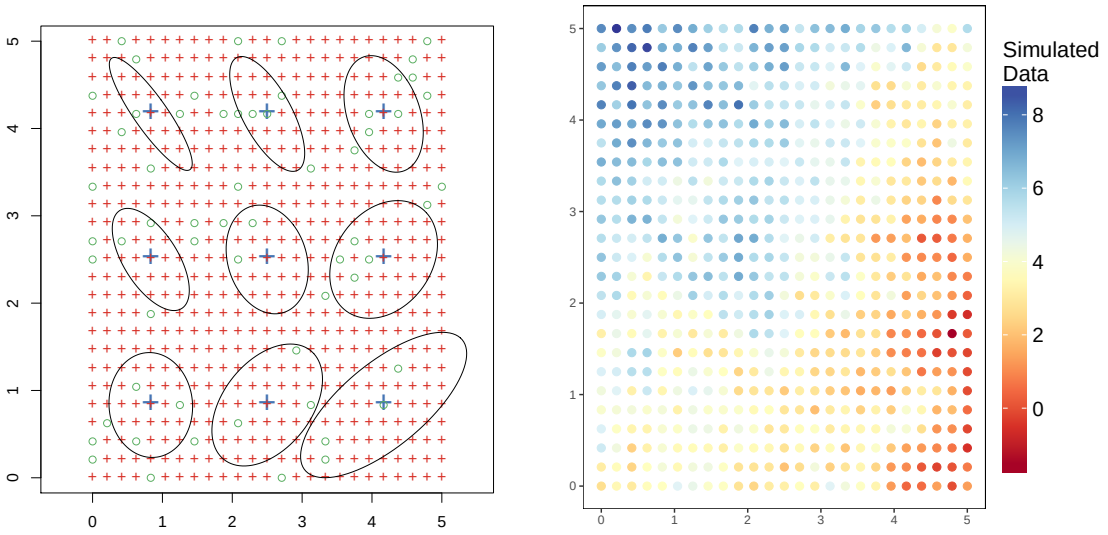

Figure 1: Left: true mixture component ellipses with observation locations (red) and validation locations (green). Right: simulated data.

6. Example 1: Simulated data

As a simple illustration, the nonstationary model will be fit to an artificial data set simulated from the model. The data lie on a 25×25 grid (so that $n = 25^2 = 625$), and there are $K = 9$ mixture component locations with corresponding mixture component ellipses as given in Figure 1. Only the kernel matrices are allowed to vary spatially, an exponential correlation structure is used, and the mean structure contains the main effects of both coordinates. The true parameter values are $\tau^2 = 0.1$, $\sigma^2 = 1$, $\beta_0 = 4$, $\beta_1 = -0.5$ (x-coordinate coefficient), $\beta_2 = 0.5$ (y-coordinate coefficient), and $\lambda_w = 2$. A total of $m = 60$ of the simulated data points are used as a validation sample. Figure 1 also provides the simulated data along with the validation locations.

The `simdata` object includes `sim.locations`, the simulated data locations; `mc.locations`, the mixture component locations; `mc.kernels`, the true mixture component kernel matrices; `sim.data`, the simulated data; and `holdout.index`, a vector of the 60 randomly sampled validation location indices. Note that there are actually ten independent and identically distributed replicates of the data contained in the ten columns of `sim.data`; in what follows the first column of data will be used.

Figure 1 can be created with

```
R> plot( simdata$mc.locations, pch = "+", asp = 1.25, xlab = "",
+    ylab = "", xlim=c( -0.5, 5.5 ), cex = 2, col = "#4575b4", main = " " )
R> points( simdata$sim.locations[ -simdata$holdout.index, ], col = "#d73027",
+    pch = "+" )
R> points( simdata$sim.locations[ simdata$holdout.index, ], col = "#5aae61" )
R> for( i in 1:dim( simdata$mc.locations )[1] ){
+    lines( ellipse( simdata$mc.kernels[ , , i ],
+    centre = simdata$mc.locations[ i, ], level = 0.5 ) )
```

```
R> ggplot( data.frame( simdata = simdata$sim.data[ , 1 ],
+    xcoord = simdata$sim.locations[ , 1 ],
+    ycoord = simdata$sim.locations[ , 2 ] ),
+    aes( x = xcoord, y = ycoord, color = simdata ) ) +
+    coord_fixed( ratio = 1.25 ) + geom_point( size = 2.5 ) +
+    xlab( "" ) + ylab( "" ) + name = "Simulated \nData") +
+    scale_color_gradientn( colours = brewer.pal( 11, "RdYlBu" ),
+    theme( panel.background = element_rect( fill = "white" ),
+    panel.border = element_rect( colour = "black", fill = NA ),
+    panel.grid = element_blank(), legend.key.height = unit( 2, "cm" ),
+    legend.title = element_text( size = 16 ),
+    legend.text = element_text( size = 15 ) )
```

(Note: this code requires the **ggplot2** (Wickham 2009), **RColorBrewer** (Neuwirth 2014), and **ellipse** (Murdoch and Chow 2013) packages in R.)

6.1. Selection of fixed components in the model

In order to fit the nonstationary model, the user must specify three components: (1) the number and placement of mixture component locations, (2) the fitting radius r, and (3) the tuning parameter λ_w. For this simulated data set, the true number and placement of the mixture component locations as well as the tuning parameter are known, and these true values will be used. However, even for the simulated data, an optimal fit radius is *not* known. When the tuning parameter is unknown the package provides a default value, but this may or may not be the optimal choice.

Given that the nonstationary model can be fit relatively quickly for given values of r, a recommended strategy is to fit the model many times for a variety of different values (same for the mixture component grid and λ_w, when these are unknown) and choose the final values based on which combination yields the "best" results. Of course, there is a trade-off between choices of the mixture component grid and fit radius: for a particular grid, the radius should be chosen such that a reasonable number of data points are used to fit each locally stationary model. It is here that the `mc_N` function may be helpful, because the user can quickly determine how many locations fall within the fitting radius of each grid point. For example, using the simulated data:

```
R> mc_N( coords = simdata$sim.locations[ -simdata$holdout.index, ],
+    mc.locations = simdata$mc.locations, fit.radius = 2 )
[1] 164 212 165 209 269 212 153 204 157
```

That is, using the true mixture component grid and a fit radius of $r = 2$, the local models will be fit using anywhere from 157 to 269 data points.

For the simulated data in this application, the "best" fit was chosen based on the quality of the anisotropy parameter estimates relative to the true values; in other cases not involving simulated data, "best" might be defined in terms of the cross-validation criteria (see, e.g., Section 7). The values used for these three components in the final fit (Section 6.2) were the true mixture component grid (with $K = 9$ evenly spaced locations over the domain), a fit radius of $r = 2.3$ units, and the true tuning parameter value of $\lambda_w = 2$.

	True value	**Stationary model**	**Nonstationary model**
β_0 (intercept)	4	4.039	3.905
β_1 (x-coordinate)	-0.5	-0.695	-0.678
β_2 (y-coordinate)	0.5	0.741	0.770
τ^2 (nugget)	0.1	0.086	0.107
σ^2 (variance)	1	1.038	0.958
CRPS	–	-0.424	-0.415
MSPE	–	0.546	0.524

Table 1: Parameter estimates from the simulated data, comparing the stationary and non-stationary models.

6.2. Final model fitting

Using the final choices of the fixed model components, the nonstationary model can be fit to the non-hold-out data (and results summarized) by

```
R> NSfit.model <- NSconvo_fit(
+    coords = simdata$sim.locations[ -simdata$holdout.index, ],
+    data = simdata$sim.data[ -simdata$holdout.index, 1 ],
+    cov.model = "exponential", fit.radius = 2.3, lambda.w = 2,
+    mc.locations = simdata$mc.locations,
+    mean.model = simdata$sim.data[ -simdata$holdout.index, 1 ]
+    ~ simdata$sim.locations[ -simdata$holdout.index, 1 ]
+    + simdata$sim.locations[ -simdata$holdout.index, 2 ] )
R> summary( NSfit.model )
```

Similarly, the anisotropic model can be fit to the non-hold-out data (and results summarized) by

```
R> anisofit.model <- Aniso_fit(
+    coords = simdata$sim.locations[ -simdata$holdout.index, ],
+    data = simdata$sim.data[ -simdata$holdout.index, 1 ],
+    cov.model = "exponential",
+    mean.model = simdata$sim.data[ -simdata$holdout.index, 1 ]
+    ~ simdata$sim.locations[ -simdata$holdout.index, 1 ]
+    + simdata$sim.locations[ -simdata$holdout.index, 2 ] )
R> summary( anisofit.model )
```

Predictions for the validation locations under each model can be calculated by

```
R> pred.NS <- predict( NSfit.model,
+    pred.coords = simdata$sim.locations[ simdata$holdout.index, ],
+    pred.covariates = simdata$sim.locations[ simdata$holdout.index, ] )
R> pred.S <- predict( anisofit.model,
+    pred.coords = simdata$sim.locations[ simdata$holdout.index, ],
```

```
+    pred.covariates = simdata$sim.locations[ simdata$holdout.index, ] )
```

(note: here, `predict` is an alias for `predict.NSconvo` and `predict.Aniso` for objects of these classes; similar syntax will be used throughout) after which the evaluation criteria can be calculated by

```
R> evaluate_CV( holdout.data = simdata$sim.data[ simdata$holdout.index ],
+    pred.mean = pred.NS$pred.means, pred.SDs = pred.NS$pred.SDs )
R> evaluate_CV( holdout.data = simdata$sim.data[ simdata$holdout.index ],
+    pred.mean = pred.S$pred.means, pred.SDs = pred.S$pred.SDs )
```

A summary of the parameter estimates and cross-validation results (MSPE and CRPS only) from each fitted model is provided in Table 1.

Calculating predictions on a finer resolution can be done as follows:

```
R> grid.x <- seq( from = 0, to = 5, by = 0.05 )
R> grid.y <- seq( from = 0, to = 5, by = 0.05 )
R> grid.locations <- expand.grid( grid.x, grid.y )
R> pred.locs <- matrix( c( grid.locations[ , 1 ], grid.locations[ , 2 ] ),
+    ncol = 2, byrow = F )
R> pred.NS.all <- predict( NSfit.model, pred.coords = pred.locs,
+    pred.covariates = pred.locs )
R> pred.S.all <- predict( anisofit.model, pred.coords = pred.locs,
+    pred.covariates = pred.locs )
```

Interpolated maps for the nonstationary model can be created with **ggplot2** (Wickham 2009):

```
R> ggplot( data.frame( preds = pred.NS.all$pred.means,
+    xcoord = pred.locs[,1],  ycoord = pred.locs[,2] ),
+    aes( x = xcoord, y = ycoord, color = preds ) ) +
+    coord_fixed( ratio = 1.25 ) + geom_point( size = 2.5 ) +
+    scale_color_gradientn( colours = brewer.pal( 11, "RdYlBu" ),
+    name = "", limits = c( -1.15, 8.7 ) ) +
+    theme( panel.background = element_rect( fill = "white" ),
+    panel.border = element_rect( colour = "black", fill = NA ),
+    panel.grid = element_blank() )
R> ggplot( data.frame( preds = pred.NS.all$pred.SDs,
+    xcoord = pred.locs[,1], ycoord = pred.locs[,2] ),
+    aes( x = xcoord, y = ycoord, color = preds ) ) +
+    coord_fixed( ratio = 1.25 ) + geom_point( size = 2.5 ) +
+    scale_color_gradientn( colours = brewer.pal( 11, "RdYlBu" ),
+    name = "", limits = c( 0.38, 0.78 ) ) +
+    theme( panel.background = element_rect( fill = "white" ),
+    panel.border = element_rect( colour = "black", fill = NA ),
+    panel.grid = element_blank() )
```

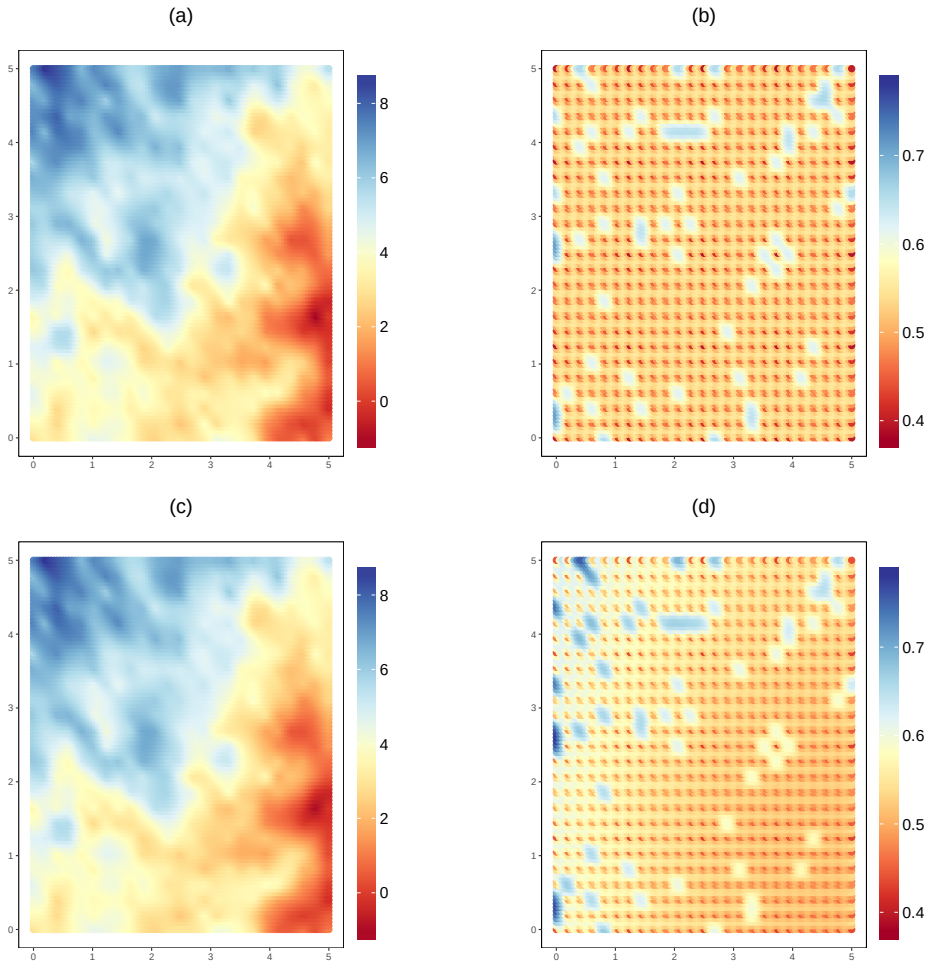

Figure 2: Predictions and prediction errors from the stationary model (a. and b.) and the nonstationary model (c. and d.).

(similarly for the stationary predictions and standard errors); these plots are provided in Figure 2. To visualize the locally-estimated anisotropy, the mixture component kernel matrices can be plotted along with the stationary anisotropy ellipse using

```
R> plot( NSfit.model, fit.radius = 2.3, xlab = "", ylab = "",
+    aniso.mat = anisofit.model$aniso.mat, asp = 1.25,
+    true.mc = simdata$mc.kernels, true.col = 1,
+    aniso.col = 4, ns.col = 2, xlim = c( -1, 6 ), ylim = c( -1, 6 ) )
```

(Note: as with `predict`, `plot` is an alias for `plot.NSconvo` for an object of class `NSconvo`; similar syntax will be used throughout.) This plot is shown in Figure 3; since the true anisotropy ellipses are known, these are also plotted. Note that the estimated nonstationary ellipses (solid red) compare quite favorably with the true ellipses (black); furthermore, the stationary ellipse (dashed blue) appears to be approximately an "average" of the spatially-varying ellipses.

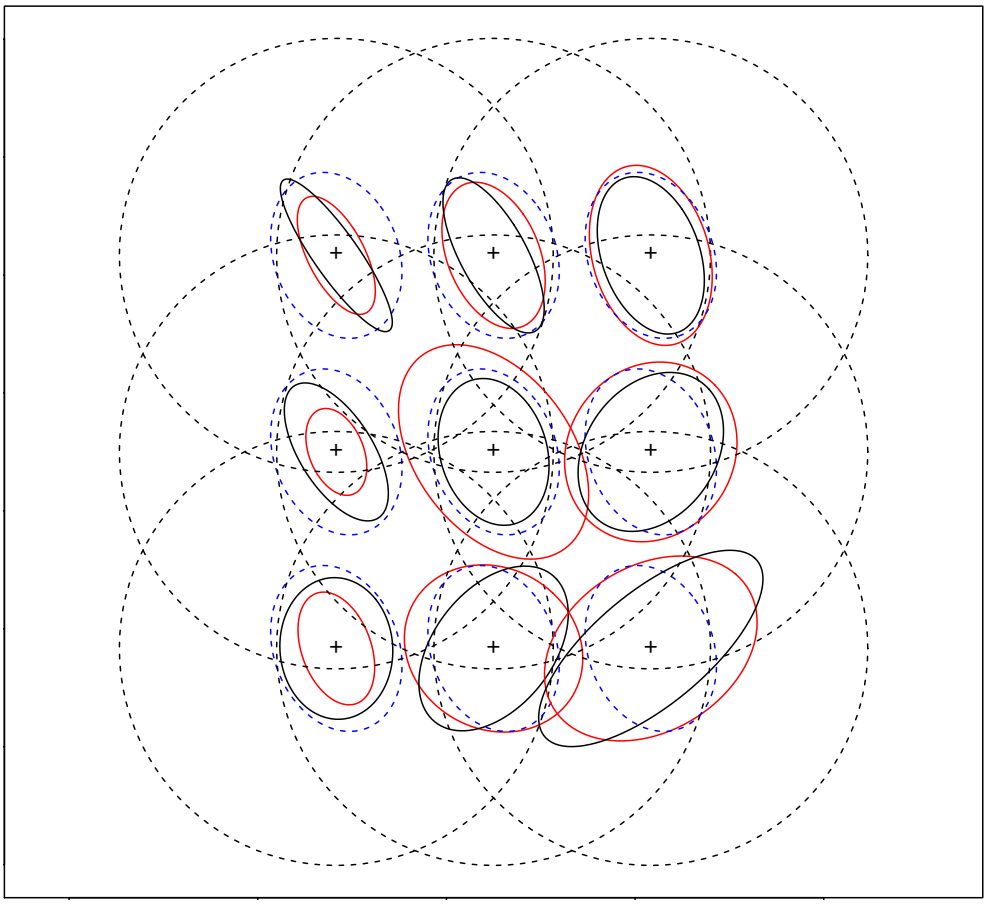

Figure 3: True mixture component ellipses (solid black) with fit radius (dashed gray), nonstationary ellipses (solid red), and the stationary ellipse (dashed blue).

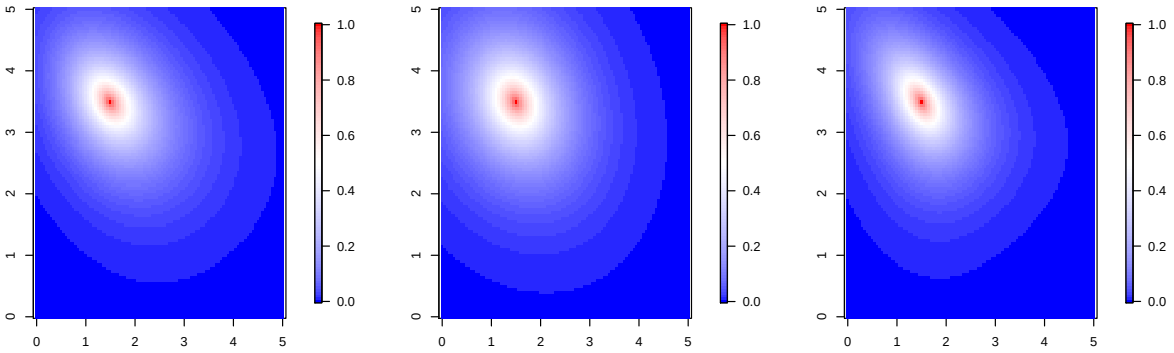

Figure 4: Estimated correlations for a reference point, showing the nonstationary (left) and stationary (center) models, as well as the true correlation (right).

As an additional visualization of the estimated nonstationarity, estimated correlation plots for a particular reference point can be obtained by

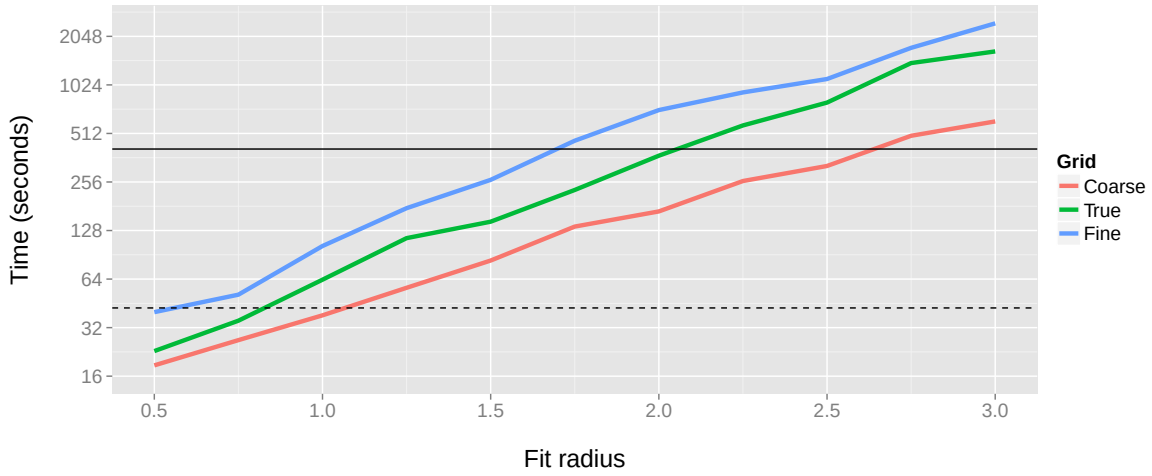

Figure 5: Computational time (in seconds) for fitting the nonstationary model to the simulated data (of size $n = 565$) across three different mixture component grids (of size $K = 4, 9, 16$, respectively) and a range of fit radii. The solid and dashed black lines represents the computational time needed to fit the stationary model using, respectively, the **convoSPAT** and **geoR** packages. Computational times given are for a Dual Quad Core Xeon 2.66GHz machine with 32GB RAM.

```
R> plot( NSfit.model, plot.ellipses = FALSE, asp = 1.25,
+    ref.loc = c( 1.5, 3.5 ), all.pred.locs = pred.locs,
+    col = diverge_hsv( 100 ) )
R> plot( anisofit.model, asp = 1.25, ref.loc = c( 1.5, 3.5 ),
+    all.pred.locs = pred.locs, col = diverge_hsv( 100 ) )
```

and are given in Figure 4 for the nonstationary and stationary models, as well as the true correlation for comparison. Note that the stationary model estimates an elliptical correlation pattern, while the nonstationary model captures the non-elliptical nature of the true correlation pattern.

Finally, note that the nonstationary model outperforms the stationary model in terms of both CRPS and MSPE (see Table 1).

6.3. Computational considerations

Selection of the size and placement of the mixture component grid and the fit radius have a major impact on the computational time. As a demonstration of this relationship, consider the computational times for a variety of choices for the mixture component grid and fit radius, summarized in Figure 5. Three mixture component grids were chosen: "coarse" (with $K = 4$), "true" (with $K = 9$), and "fine" (with $K = 16$). For each of these grids, the nonstationary model (with spatially-varying nugget and variance) was fit using fit radii ranging between $r = 0.5$ and $r = 3$. As a reference point, the average number of data points used to fit the locally stationary models increases linearly from around 20 for $r = 0.5$ to 350 for $r = 3$ (consistent across the different mixture component grids).

The linear trends on the log scale in Figure 5 are expected, as calculations for Gaussian

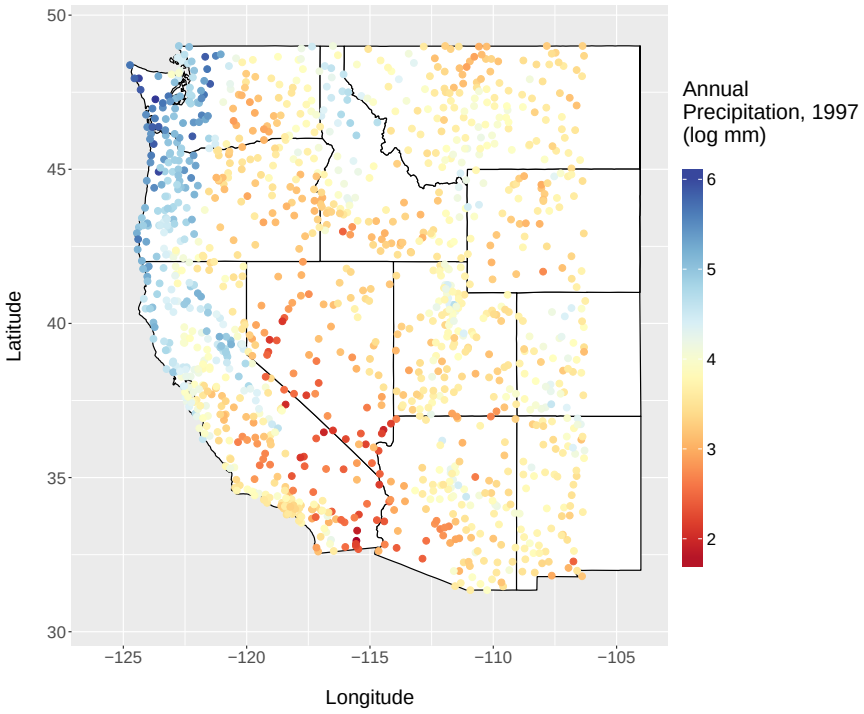

Figure 6: Annual precipitation data for 1997, in log mm.

process models are known to increase with the square of the sample size. The (approximately) vertical shifts present in moving from coarse to true to fine are also expected, as the finer grids represent fitting more local models with (approximately) the same local sample sizes.

For comparison, the time required to fit the stationary model using `Aniso_fit` is also shown on Figure 5 as a solid black line. However, as mentioned in Section 5.2, if the user is interested in fitting a stationary model to a spatial data set, a much better choice is `likfit` in **geoR** (Ribeiro Jr. and Diggle 2015). The corresponding stationary (anisotropic) model in **geoR** can be fit by

```
R> geoR.fit <- likfit( data = simdata$sim.data[ -simdata$holdout.index, 1 ],
+     coords = simdata$sim.locations[ -simdata$holdout.index, ],
+     cov.model = "exponential", ini.cov.pars = c( 1, 1 ), trend = "1st",
+     fix.psiA = FALSE, fix.psiR = FALSE, lik.method = "REML" )
```

The computational time for `likfit` is also shown on Figure 5 as a dashed black line.

7. Example 2: Annual precipitation

The nonstationary model proposed in Sections 3 and 4 is also applied to a moderately large, real data set, consisting of the total annual precipitation in the western United States for 1997. The data is available online from the National Center for Atmospheric Research (`http://www.image.ucar.edu/GSP/Data/US.monthly.met`) as part of a larger data set that includes measurements for the entire United States. For the purposes of this analysis, a subset of the

Label	Covariance model details
Stationary	Anisotropic, constant nugget and variance
NS1	SV anisotropy, constant nugget and variance
NS2	SV anisotropy and variance, constant nugget
NS3	SV anisotropy and nugget, constant variance
NS4	SV anisotropy, variance, and nugget

Table 2: A brief summary of the different models fit to the precipitation data. Note: "SV" indicates "spatially-varying."

data that includes the western United States was chosen (see Figure 6) because precipitation is smoother and more densely observed over the central and eastern United States. This subset is included in the package as an `RData` (named `USprecip97`) file and consists of 1270 observations. For illustration purposes, we work with non-projected longitude and latitude, but note that due to the size of the study area this may be inappropriate.

7.1. Spatial model summaries

A total of five spatial models were fit to this dataset, the details of which are summarized in Table 2: the stationary model and four nonstationary models. All models were assigned the same mean structure, which included the main effects of longitude and latitude as well as an intercept. Additionally, all models were assigned the same underlying correlation structure, chosen to be exponential (i.e., the Matérn class with fixed $\kappa = 0.5$). Twenty percent of the observations ($m = 254$) were held out as a validation data set in order to evaluate each model, leaving $n = 1016$ observations as a training data set. The stationary model was fit using the `Aniso_fit` function and the nonstationary models were fit using the `NSconvo_fit` function; the four nonstationary models represent all combinations of the `ns.nugget` and `ns.variance` options.

In addition to selecting which of the nugget variance and/or process variance should be spatially-varying, recall that using this model requires specifying (1) a mixture component grid, (2) the tuning parameter for the weight function, and (3) the fitting radius r. For each of the nonstationary models, the model was actually fit multiple times using two different mixture component grids ("coarse", with $K = 15$, and "fine", with $K = 22$), four different values of the tuning parameter ($\lambda_w = 1.00, 2.67, 4.33, 6.00$), and six different fit radii (for the coarse grid, $r = 2.5, 3.2, 3.9, 4.6, 5.3, 6.0$; for the fine grid, $r = 2.75, 3.20, 3.65, 4.10, 4.55, 5.00$).

Of the $2 \times 4 \times 6 = 48$ total models fit for each of the four nonstationary models, the best model was selected based on maximizing the CRPS criteria (for the test data using the log precipitation); the best models are summarized in Table 3. For each model, the coarse grid with a fit radius of $r = 3.9$ performed best, and the largest tuning parameter was preferred for all but model NS1 (however, the CRPS for NS1 with $\lambda_w = 6$ was -0.13714, which is only slightly smaller than the CRPS for $\lambda_w = 4.33$, which was -0.13709). Table 3 also provides the computational run time for each model. Parameter estimates for each of the best models are summarized in Table 4.

While the gains are modest, all of the nonstationary models outperform the stationary model in terms of both MSPE and CRPS; the best model in terms of both MSPE and CRPS is NS1.

Model	**Grid size**	λ_w	r	**CRPS**	**MSPE**	**Comp. time**
Stationary	–	–	–	-0.1424	0.0666	85.94 min.
NS1	15 (coarse grid)	4.33	3.9	-0.1371	0.0610	17.02 min.
NS2	15 (coarse grid)	6.0	3.9	-0.1387	0.0619	16.31 min.
NS3	15 (coarse grid)	6.0	3.9	-0.1377	0.0619	16.55 min.
NS4	15 (coarse grid)	6.0	3.9	-0.1386	0.0623	9.49 min.

Table 3: Model details and evaluation for the best model of each type fit to the precipitation data, selected based on maximizing CRPS. The computational time given is for a Dual Quad Core Xeon 2.66GHz machine with 32GB RAM.

Parameter	**S**	**NS1**	**NS2**	**NS3**	**NS4**
β_0 (int.)	-5.315	-2.802	-1.138	-2.916	-1.167
β_1 (x-coord)	-0.067	-0.037	-0.028	-0.039	-0.029
β_2 (y-coord)	0.034	0.055	0.0038	0.053	0.038
λ_1	2.772	SV	SV	SV	SV
λ_2	6.385	SV	SV	SV	SV
η	0.336	SV	SV	SV	SV
τ^2	0.011	0.005	0.013	SV	SV
σ^2	0.379	0.289	SV	0.275	SV

Table 4: Parameter estimates for the five best spatial models fit to the precipitation data set, indicating which parameters are spatially-varying (SV). Recall that with an exponential correlation structure, the smoothness κ is fixed to be 0.5 for all models.

The interpolated prediction maps and corresponding standard errors for the stationary model and NS1 are given in Figure 7. Note that the nonstationary model estimates the prediction errors much more flexibly, with larger errors in the far southwest (southern California, southern Nevada, and western Arizona) and smaller errors in the northern portion of the domain. The stationary model prediction errors are much more homogeneous, being highly dependent on the proximity of neighboring observations.

7.2. Visualizations of nonstationarity

Graphical summaries can be made to visualize the spatially-varying parameter estimates not explicitly provided in Table 4. The fully nonstationary model NS4 will be used for the following visualizations, as all of its variance/covariance parameters are spatially-varying.

First, consider the locally-estimated anisotropy ellipses for each of the mixture component locations, shown by the solid red ellipses in Figure 8. This plot also contains the corresponding ellipse from the stationary model (dashed blue), which is constant over space, and the neighborhood size (dotted gray) defined by the fit radius r. It is quite clear that precipitation in the western United States displays nonstationary behavior with respect to the anisotropy ellipses, as the local estimates are highly variable across the spatial domain.

Next, consider the estimates of the spatially-varying nugget variance (τ^2) and process variance

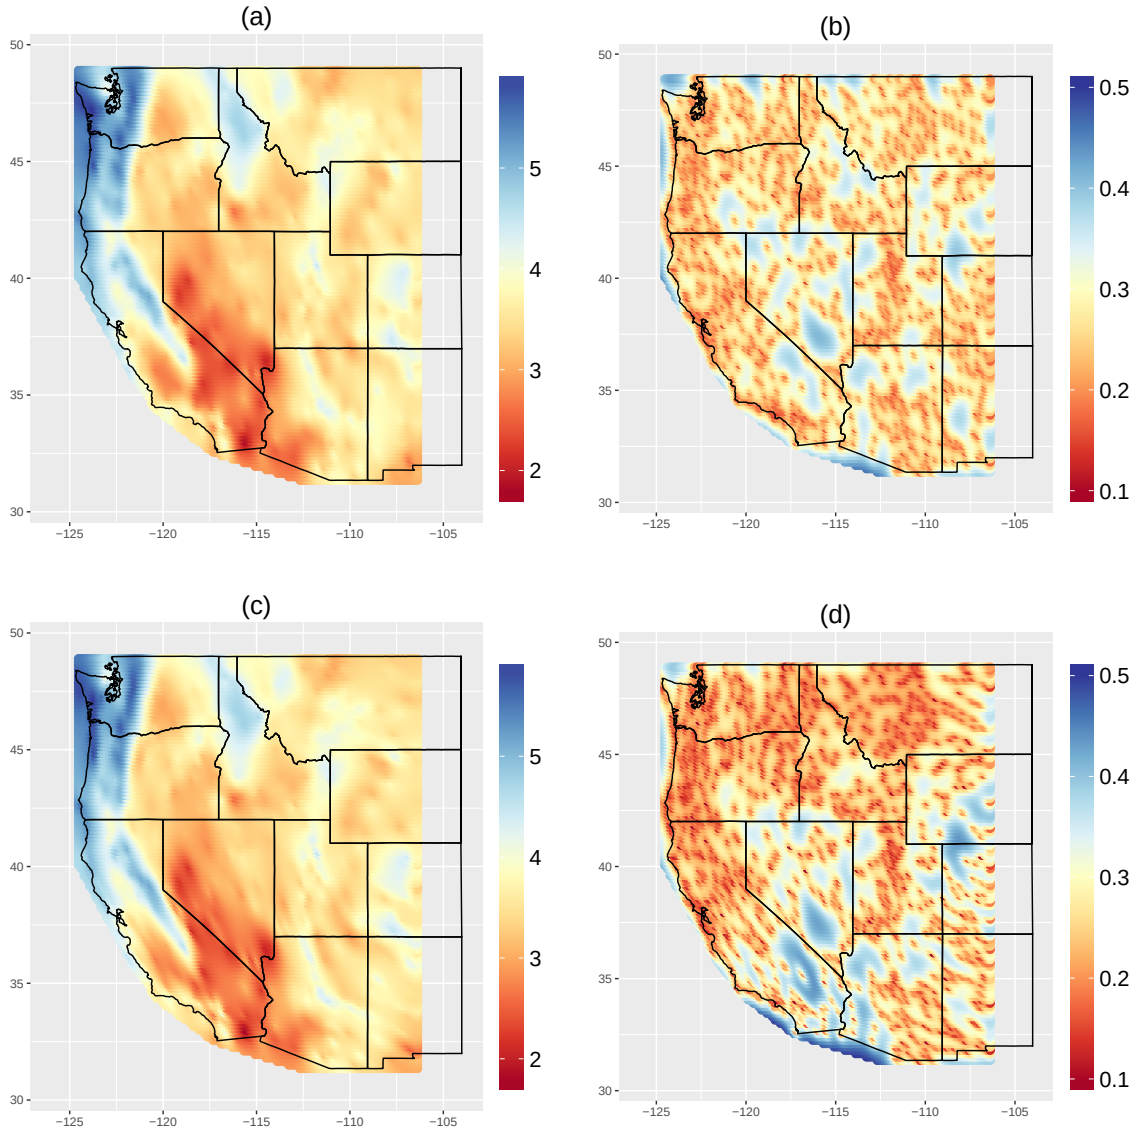

Figure 7: Predictions and prediction standard errors for the stationary model (plots (a) and (b)) and the nonstationary model NS1 (plots (c) and (d)).

(σ^2) over the spatial region for NS4, shown in Figure 9. The nugget variance is highly variable across the spatial domain, which may reflect variability in the monitoring devices. The process variance is largest along the coastal United States, which is intuitive based on its proximity to the ocean and highly variable topography. In general, areas with larger process variance correspond to smaller nugget variance, although this is not true everywhere (e.g., central Montana).

Finally, estimated correlation plots for three reference points are given in Figure 10. This plot nicely illustrates the fact that the nonstationary model allows the spatial dependence structure to change over space, while the stationary model estimates a constant correlation structure.

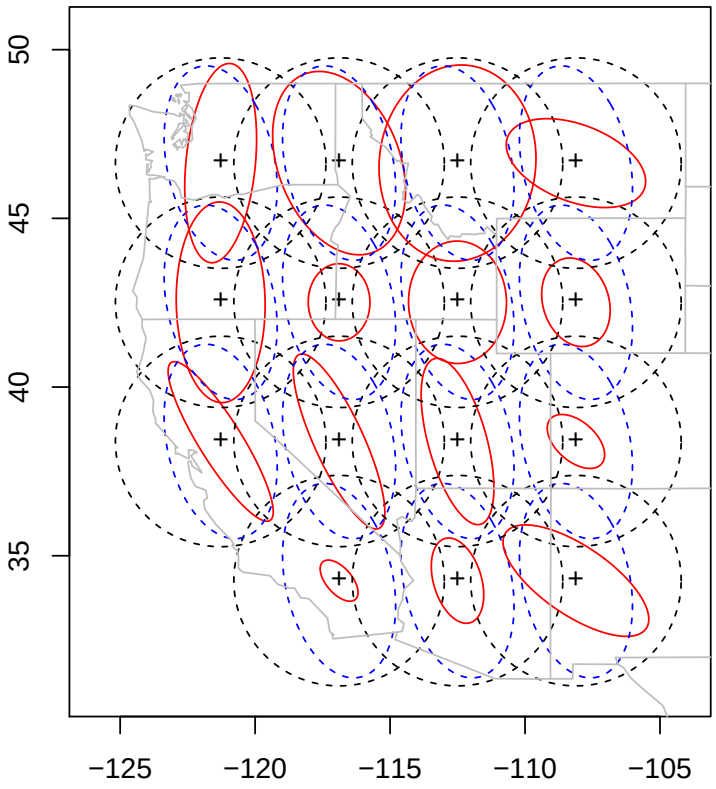

Figure 8: Estimated mixture component ellipses for the nonstationary model (red), the stationary (anisotropic) model (blue), and the estimation region (dashed black).

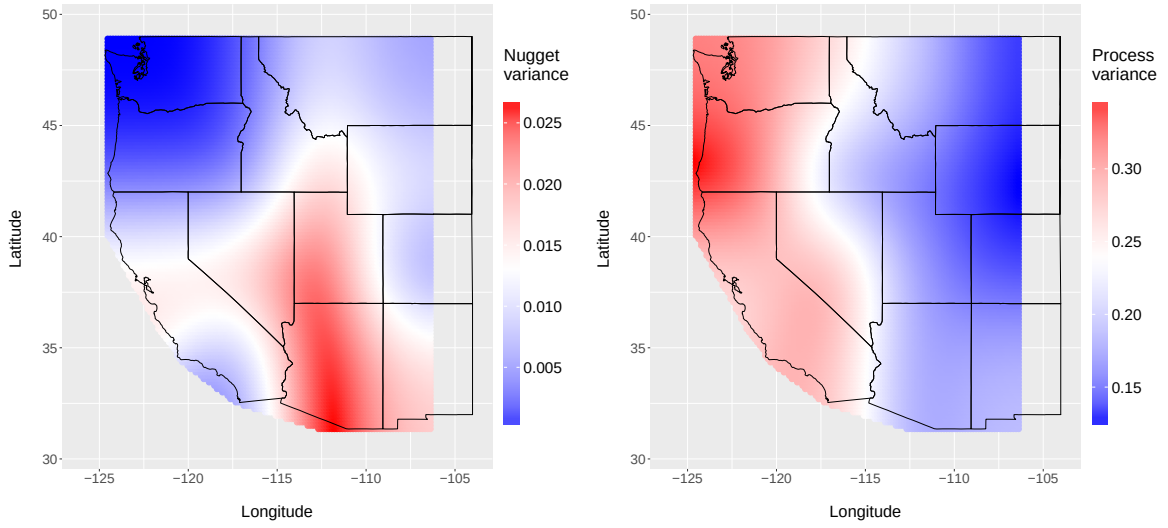

Figure 9: Plots of the estimated spatially-varying process variance σ^2 (right) and nugget variance τ^2 (left) for model NS4.

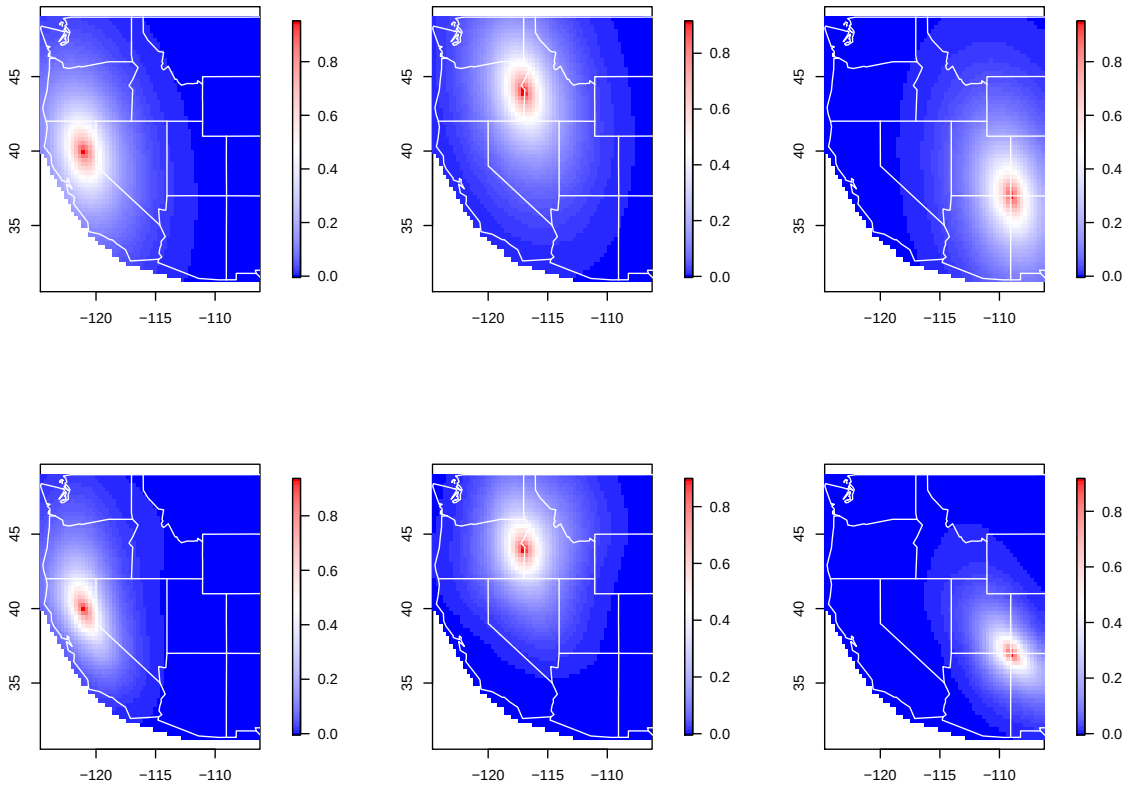

Figure 10: Correlation plots for three reference points, comparing the stationary model S (top) and nonstationary model NS4 (bottom).

In addition to allowing the orientation of the anisotropy ellipses to change over space (as seen in Figure 8), the nonstationary model allows for non-elliptical correlation patterns.

8. Discussion

In this paper, we have presented a nonstationary spatial Gaussian process model that is highly flexible yet amenable to computationally efficient inference, as shown through its implementation in the new **convoSPAT** package for R. In fact, for a moderately large real data set, fitting the nonstationary model is significantly faster than fitting the stationary model (at least using `Aniso_fit`; see Table 3). The model also allows for visualization of the estimated nonstationarity, for both the spatially-varying variance parameters and correlation structure.

The tradeoff for the computational tractability of this model is that uncertainty in the locally-estimated parameters is not accounted for in global estimation and, more seriously, this uncertainty is not quantified in the parameter estimation. Furthermore, the model is not completely pre-packaged, in that the user must specify the mixture component grid, fit radius, and tuning parameter, and the resulting parameter estimates and predictions may be sensitive to these

choices. Since uncertainty in parameter estimates is not provided, it is difficult to determine if a nonstationary model is needed or if a stationary model would be sufficient. However, the primary goal of this model is to provide a "quick and dirty" way to fit a nonstationary Gaussian process model to spatial data, allowing new nonstationary methods to be compared with another model that is also nonstationary, instead of simply a stationary model. And, since the model can be fit very quickly, a practitioner can fit the model for many choices of the mixture component grid, fit radius, and tuning parameter (similar to the strategy in Section 7.1) and use the provided evaluation criteria (or other criteria) to choose a final model.

Finally, while this model provides a fast approach for modeling nonstationary spatial Gaussian processes, note that the same model could gain an additional degree of computational efficiency by parallelizing the LLE component of the model. Currently, the local parameters are estimated sequentially; however, the optimization for each mixture component location does not depend on the optimization for the other mixture component locations. Thus, a more efficient algorithm might estimate all of these components simultaneously, greatly reducing the overall computation time even further. A future version of **convoSPAT** will explore this possibility.

References

Adler RJ (1981). *The Geometry of Random Fields*. John Wiley & Sons.

Anderes EB, Stein ML (2011). "Local Likelihood Estimation for Nonstationary Random Fields." *Journal of Multivariate Analysis*, **102**(3), 506 – 520. ISSN 0047-259X. doi:http://dx.doi.org/10.1016/j.jmva.2010.10.010. URL http://www.sciencedirect.com/science/article/pii/S0047259X10002162.

Barry RP, Ver Hoef JM (1996). "Blackbox Kriging: Spatial Prediction Without Specifying Variogram Models." *Journal of Agricultural, Biological, and Environmental Statistics*, **1**(3), 297–322.

Bivand RS, Pebesma E, Gomez-Rubio V (2013). *Applied Spatial Data Analysis with R*. Second edition. Springer-Verlag, N.Y. URL http://www.asdar-book.org/.

Bochner S (1959). *Lectures on Fourier integrals. No. 42*. Princeton University Press.

Calder CA (2008). "A Dynamic Process Convolution Approach to Modeling Ambient Particulate Matter Concentrations." *Environmetrics*, **19**(1), 39–48. ISSN 1099-095X. doi:10.1002/env.852. URL http://dx.doi.org/10.1002/env.852.

Damian D, Sampson PD, Guttorp P (2001). "Bayesian Estimation of Semi-Parametric Non-Stationary Spatial Covariance Structures." *Environmetrics*, **12**(2), 161–178. ISSN 1099-095X. doi:10.1002/1099-095X(200103)12:2<161::AID-ENV452>3.0.CO;2-G. URL http://dx.doi.org/10.1002/1099-095X(200103)12:2<161::AID-ENV452>3.0.CO;2-G.

D'Orazio M (2016). ***StatMatch**: Statistical Matching*. R package version 1.2.4, URL https://CRAN.R-project.org/package=StatMatch.

Fuentes M (2001). "A High Frequency Kriging Approach for Non-Stationary Environmental Processes." *Environmetrics*, **12**(5), 469–483. ISSN 1099-095X. doi:10.1002/env.473. URL http://dx.doi.org/10.1002/env.473.

Fuentes M (2002). "Spectral Methods for Nonstationary Spatial Processes." *Biometrika*, **89**(1), 197–210. doi:10.1093/biomet/89.1.197. http://biomet.oxfordjournals.org/content/89/1/197.full.pdf+html, URL http://biomet.oxfordjournals.org/content/89/1/197.abstract.

Fuglstad GA, Lindgren F, Simpson D, Rue H (2015). "Exploring a New Class of Non-stationary Spatial Gaussian Random Fields with Varying Local Anisotropy." *Statistica Sinica*, **25**, 115–133. doi:10.5705/ss.2013.106w.

Gneiting T, Raftery AE (2007). "Strictly Proper Scoring Rules, Prediction, and Estimation." *Journal of the American Statistical Association*, **102**(477), 359–378. doi:10.1198/016214506000001437. http://dx.doi.org/10.1198/016214506000001437, URL http://dx.doi.org/10.1198/016214506000001437.

Higdon D (1998). "A Process-Convolution Approach to Modeling Temperatures in the North Atlantic Ocean." *Environmental and Ecological Statistics*, **5**(2), 173–190. ISSN 1352-8505. doi:10.1023/A:1009666805688. URL http://dx.doi.org/10.1023/A3A1009666805688.

Higdon D (2002). "Space and Space-Time Modeling Using Process Convolutions." In C Anderson, V Barnett, P Chatwin, A El-Shaarawi (eds.), *Quantitative Methods for Current Environmental Issues*, pp. 37–56. Springer-Verlag London. ISBN 978-1-4471-1171-9. doi:10.1007/978-1-4471-0657-9_2. URL http://dx.doi.org/10.1007/978-1-4471-0657-9_2.

Hoef JMV, Barry RP (1998). "Constructing and Fitting Models for Cokriging and Multivariable Spatial Prediction." *Journal of Statistical Planning and Inference*, **69**(2), 275 – 294. ISSN 0378-3758. doi:http://dx.doi.org/10.1016/S0378-3758(97)00162-6. URL http://www.sciencedirect.com/science/article/pii/S0378375897001626.

Ickstadt K, Wolpert RL (1998). "Spatial Regression for Marked Point Processes." *Bayesian Statistics*, **6**.

Ingebrigtsen R, Lindgren F, Steinsland I (2014). "Spatial Models with Explanatory Variables in the Dependence Structure." *Spatial Statistics*, **8**(0), 20 – 38. ISSN 2211-6753. doi:http://dx.doi.org/10.1016/j.spasta.2013.06.002. URL http://www.sciencedirect.com/science/article/pii/S2211675313000377.

Katzfuss M (2013). "Bayesian Nonstationary Spatial Modeling for Very Large Datasets." *Environmetrics*, **24**(3), 189–200. ISSN 1099-095X. doi:10.1002/env.2200. URL http://dx.doi.org/10.1002/env.2200.

Kitanidis PK (1983). "Statistical Estimation of Polynomial Generalized Covariance Functions and Hydrologic Applications." *Water Resources Research*, **19**(4), 909–921. ISSN 1944-7973. doi:10.1029/WR019i004p00909. URL http://dx.doi.org/10.1029/WR019i004p00909.

Kleiber W, Nychka D (2012). "Nonstationary Modeling for Multivariate Spatial Processes." *Journal of Multivariate Analysis*, **112**(0), 76 – 91. ISSN 0047-259X. doi:http://dx.doi.org/10.1016/j.jmva.2012.05.011. URL http://www.sciencedirect.com/science/article/pii/S0047259X12001376.

Lemon J (2006). "**plotrix**: A Package in the Red Light District of R." *R-News*, **6**(4), 8–12.

Lindgren F, Rue H (2015). "Bayesian Spatial Modeling with R-**INLA**." *Journal of Statistical Software*, **63**(19), 1–25. doi:10.1111/j.1467-9868.2011.00777.x. URL http://www.jstatsoft.org/v63/i19/.

Lindgren F, Rue H, Lindstrom J (2011). "An Explicit Link Between Gaussian Fields and Gaussian Markov Random Fields: The Stochastic Partial Differential Equation Approach." *Journal of the Royal Statistical Society B*, **73**(4), 423–498. ISSN 1467-9868. doi:10.1111/j.1467-9868.2011.00777.x. URL http://dx.doi.org/10.1111/j.1467-9868.2011.00777.x.

Murdoch D, Chow ED (2013). ***ellipse**: Functions for Drawing Ellipses and Ellipse-Like Confidence Regions*. R package version 0.3-8, URL http://CRAN.R-project.org/package=ellipse.

Neuwirth E (2014). ***RColorBrewer**: ColorBrewer Palettes*. R package version 1.1-2, URL http://CRAN.R-project.org/package=RColorBrewer.

Nychka D, Furrer R, Sain S (2014). ***fields**: Tools for Spatial Data*. R package version 7.1, URL http://CRAN.R-project.org/package=fields.

Paciorek CJ (2003). *Nonstationary Gaussian Processes for Regression and Spatial Modeling*. Ph.D. thesis, Carnegie Mellon University.

Paciorek CJ, Schervish MJ (2006). "Spatial Modeling Using a New Class of Nonstationary Covariance Functions." *Environmetrics*, **17**, 483–506.

Patterson HD, Thompson R (1971). "Recovery of Inter-Block Information When Block Sizes Are Unequal." *Biometrika*, **58**(3), 545–554. doi:10.1093/biomet/58.3.545. http://biomet.oxfordjournals.org/content/58/3/545.full.pdf+html, URL http://biomet.oxfordjournals.org/content/58/3/545.abstract.

Patterson HD, Thompson R (1974). "Maximum Likelihood Estimation of Components of Variance." Proc. Eighth International Biochem. Conf.

Pebesma E, Bivand R (2005). "Classes and Methods for Spatial Data in R." *R-News*, **5**(2). URL http://cran.r-project.org/doc/Rnews/.

Reich BJ, Eidsvik J, Guindani M, Nail AJ, Schmidt AM (2011). "A Class of Covariate-Dependent Spatiotemporal Covariance Functions for the Analysis of Daily Ozone Concentration." *The Annals of Applied Statistics*, **5**(4), 2425–2447. doi:10.1214/11-AOAS482. URL http://dx.doi.org/10.1214/11-AOAS482.

Ribeiro Jr PJ, Diggle PJ (2015). ***geoR**: Analysis of Geostatistical Data*. R package version 1.7-5.1, URL http://CRAN.R-project.org/package=geoR.

Risser MD, Calder CA (2015). "Regression-Based Covariance Functions for Nonstationary Spatial Modeling." *Environmetrics*, **26**(4), 284–297. ISSN 1099-095X. doi:10.1002/env.2336. URL http://dx.doi.org/10.1002/env.2336.

Sampson PD, Guttorp P (1992). "Nonparametric Estimation of Nonstationary Spatial Covariance Structure." *Journal of the American Statistical Association*, **87**(417), 108–119. doi:10.1080/01621459.1992.10475181. http://www.tandfonline.com/doi/pdf/10.1080/01621459.1992.10475181, URL http://www.tandfonline.com/doi/abs/10.1080/01621459.1992.10475181.

Schmidt AM, Guttorp P, O'Hagan A (2011). "Considering Covariates in the Covariance Structure of Spatial Processes." *Environmetrics*, **22**(4), 487–500. ISSN 1099-095X. doi:10.1002/env.1101. URL http://dx.doi.org/10.1002/env.1101.

Schmidt AM, O'Hagan A (2003). "Bayesian Inference for Non-Stationary Spatial Covariance Structure Via Spatial Deformations." *Journal of the Royal Statistical Society B*, **65**(3), 743–758. ISSN 1467-9868. doi:10.1111/1467-9868.00413. URL http://dx.doi.org/10.1111/1467-9868.00413.

Stein ML (2005). "Nonstationary Spatial Covariance Functions." *Unpublished technical report*.

Thiébaux HJ (1976). "Anisotropic Correlation Functions for Objective Analysis." *Monthly Weather Review*, **104**, 994–1002.

Thiébaux HJ, Pedder MA (1987). *Spatial Objective Analysis: With Applications in Atmospheric Science*. Academic Press.

Tibshirani R, Hastie T (1987). "Local Likelihood Estimation." *Journal of the American Statistical Association*, **82**(398), 559–567. ISSN 01621459. URL http://www.jstor.org/stable/2289465.

Venables WN, Ripley BD (2002). *Modern Applied Statistics With S*. Fourth edition. Springer-Verlag, New York. ISBN 0-387-95457-0, URL http://www.stats.ox.ac.uk/pub/MASS4.

Ver Hoef JM, Cressie N, Barry RP (2004). "Flexible Spatial Models for Kriging and Cokriging Using Moving Averages and the Fast Fourier Transform (FFT)." *Journal of Computational and Graphical Statistics*, **13**(2), 265–282. doi:10.1198/1061860043498. http://dx.doi.org/10.1198/1061860043498, URL http://dx.doi.org/10.1198/1061860043498.

Vianna Neto JH, Schmidt AM, Guttorp P (2014). "Accounting for Spatially Varying Directional Effects in Spatial Covariance Structures." *Journal of the Royal Statistical Society C*, **63**(1), 103–122. ISSN 1467-9876. doi:10.1111/rssc.12027. URL http://dx.doi.org/10.1111/rssc.12027.

Wickham H (2009). ***ggplot2**: Elegant Graphics for Data Analysis*. Springer-Verlag New York. ISBN 978-0-387-98140-6. URL http://had.co.nz/ggplot2/book.

Affiliation:

Catherine A. Calder
Department of Statistics

The Ohio State University
1958 Neil Avenue
Columbus, OH, USA 43210
E-mail: calder@stat.osu.edu